

Weak randomness in graphons and theons

Leonardo N. Coregliano* Maryanthe Malliaris†

July 22, 2024

Abstract

Call a hereditary family \mathcal{F} of graphs strongly persistent if there exists a graphon W such that in all subgraphons W' of W , \mathcal{F} is precisely the class of finite graphs that have positive density in W' . Our first result is a complete characterization of the hereditary families of graphs that are strongly persistent as precisely those that are closed under substitutions.

We call graphons with the self-similarity property above weakly random. A hereditary family \mathcal{F} is said to have the weakly random Erdős–Hajnal property (WR) if every graphon that is a limit of graphs in \mathcal{F} has a weakly random subgraphon. Among families of graphs that are closed under substitutions, we completely characterize the families that belong to WR as those with “few” prime graphs.

We also extend some of the results above to structures in finite relational languages by using the theory of theons.

Keywords: Graph limit, theon, quasirandomness.

1 Introduction

The theory of graph quasirandomness implies that quasirandom graphons are the only graphons W with the self-similarity property that densities of finite graphs are invariant across subgraphons of W (see [Tho87, CGW89] for graph quasirandomness and [Lov12] for graphons). An interesting weakening of this property, which we will motivate further below, is to require only that the family \mathcal{F} of finite graphs that have positive density is invariant across subgraphons of W . We call graphons with this property *weakly random*. It is natural to ask which families \mathcal{F} can be realized in this way in some weakly random W . Since all constant graphons are quasirandom, thus also weakly random, three such families are the cliques, the anti-cliques and the family of all finite graphs. However, there are other families that can be realized in this way such as the family \mathcal{F}_{C_4} of all cographs, that is, graphs such that every induced subgraph of size at least 2 can be partitioned into two non-trivial

*Institute for Advanced Study, lenacore@ias.edu. This material is based upon work supported by a grant from the Institute for Advanced Study School of Mathematics.

†University of Chicago, mem@math.uchicago.edu. Research partially supported by NSF-BSF 2051825.

parts that are either complete to each other, or empty to each other. Alternatively, \mathcal{F}_{C_4} is precisely the family of finite graphs that are induced subgraphs of some recursive blow-up of the 4-cycle. Strong persistence of \mathcal{F}_{C_4} is seen since the limit W_{C_4} of the balanced recursive blow-ups of the 4-cycle is weakly random and realizes the family \mathcal{F}_{C_4} .

The work of this paper is to show that this notion of weak randomness supports a rich structure theory and provides an illuminating way of studying hereditary classes of graphs based on properties of their limit objects. Before stating our main results, let us further motivate why the study of weak randomness is both natural and tractable, which begins by asking what is special about large cliques and anti-cliques.

Recall that the Erdős–Hajnal Conjecture [EH89] says that for any proper hereditary class of graphs, there exists a constant $c > 0$ such that any graph of size n in this class either has a clique or an anti-clique of size n^c ; we will refer to this property of a hereditary class as the Erdős–Hajnal property and abbreviate it as **EHP** (see also [Chu14] for a survey). In [CM22], we studied a natural variant of this question in the presence of convergence, called the approximate Erdős–Hajnal property (**AEHP**), in which we allow for a negligible amount of non-edges in the almost clique or a negligible amount of edges in the almost anti-clique, but require it to be linear-sized. The framework of **AEHP** naturally lends itself to analysis via limit theory, i.e., graphons [LS06] in the graph case, or more generally, flag algebras [Raz07] and theons [CR20] in the case of universal theories in finite relational languages.

The aforementioned family \mathcal{F}_{C_4} of cographs plays a key role in some of the classical results on the usual Erdős–Hajnal Conjecture: namely, a consequence of [APS01, Theorem 1.1] is that any hereditary class that does not contain \mathcal{F}_{C_4} has **EHP**. However, this is not a characterization of **EHP** as several classes that contain \mathcal{F}_{C_4} still have **EHP**; easy examples include perfect graphs and P_3 -free graphs (i.e., disjoint unions of cliques) and hard examples include bull-free graphs [CS08] and C_5 -free graphs [CSSS23].

On the other hand, surprisingly, hereditary classes of graphs with **AEHP** can be characterized as precisely those that avoid containing \mathcal{F}_{C_4} , see [CM22, Theorem 8.10]. In what follows, it will be more convenient to think about hereditary classes of graphs as the models of a particular universal first-order theory T of graphs, so a graphon of T is simply a limit of finite models of T . This shift in language supports the model theoretic perspective of studying the theory T (i.e., a hereditary class of graphs) by studying the variation in the class of its infinite models (i.e., its graphons). In this language, a universal theory T of graphs has **AEHP** if every graphon of T has a (large) trivial subgraphon, i.e., an almost clique or an almost anti-clique, see [CM22, §7] and Definition 2.9.

In the proof of the negative side of the characterization of **AEHP** for graphs, if all cographs are models of T , then the limit W_{C_4} is a graphon of T . Looking through the lens of weak randomness, it is clear that W_{C_4} does not contain trivial subgraphons since both the edge and the non-edge must persistently have positive density in all subgraphons of W_{C_4} . Part of the characterization of **AEHP** involved showing that persistence of the edge and non-edge implies persistence of every graph in \mathcal{F}_{C_4} . Thus, we are led to ask which families arise as *persistent classes* of graphons, i.e., families \mathcal{F} of graphs that are precisely those that have positive density in all subgraphons of a given graphon W . A related notion is that of a

strongly persistent class, in which the graphon is further required to be weakly random. A priori these notions are different since a non-weakly random graphon can have finite graphs with positive density in only some of its subgraphons.

The first theorem of the present paper is to show the equivalence of strong persistence and persistence and to characterize such families as precisely those that are closed under substructures and substitutions (see Definition 3.8). This requires both understanding properties of substitutions and the construction of appropriate weakly random limits. We prove this result first for graphs (Theorem 4.4) and then a suitable generalization of it for structures in arbitrary finite relational languages (Theorem 7.9) after developing suitable extensions of the relevant concepts. The appearance of substitution in this characterization, and of the related notion of primality in what follows, is not completely unexpected as both the Erdős–Hajnal property and its approximate version behave very well under substitution (see [APS01, Chu14] for EHP and [CM22] for AEHP).

Since cliques and anti-cliques are weakly random, we can extend the picture of AEHP by defining the class WR as follows: a universal theory of graphs is in AEHP if all its graphons have trivial subgraphons and a universal theory of graphs is in WR if all its graphons have weakly random subgraphons. It is immediate that $\text{AEHP} \subseteq \text{WR}$, it is less immediate but shown in the present paper that this containment is proper and that not every universal theory is in WR. Because of the nature and simplicity of the characterization of AEHP for graphs cited above, it becomes plausible that a characterization of the richer WR class may exist.

In Theorem 5.4, we characterize theories of graphs in WR under the additional natural assumption of closure under substitution as those that have “few” prime graphs in the sense that there are no infinite antichains of prime graphs in the induced subgraph partial order, a condition we call *primally almost finite*. In one direction, we build on the analysis of persistence of Theorem 4.4 and in the other direction, the technology of recursive blow-ups plays a key role. Note that without the assumption of closure under substitutions, it is obvious that WR is no longer characterized by the primally almost finite condition as, e.g., the theory of bipartite graphs is in WR (even in AEHP) but has infinite antichains of prime graphs.

Many further questions are discussed in the concluding Section 9.

Let us point out that although [CM22] provides a good motivation for the current work, it is not a pre-requisite for the current paper and we do not rely on any of the results of [CM22] for our study of weak randomness and the class WR, except for a straightforward characterization of subgraphons and sub-objects [CM22, Lemmas 3.3 and 5.8] (see also Section 2 below). To read the current paper, it will be useful to have some familiarity with the theories of graphons and theons, but we repeat the relevant definitions and results in Section 2 to set the notation.

Now we describe the structure of the paper. In Section 2, we review the necessary preliminaries and set notation. Section 3 starts to develop the properties of substitution, primality and almost finiteness, which we will need for the rest of the paper. Section 4 is devoted to proving the persistence result for graphs, Theorem 4.4. Section 5 defines the class WR for graphs and proves the characterization under the assumption of closure under substitutions,

Theorem 5.4. In Section 6, we study how the notions of weak randomness interact with VC dimension, show that weakly random graphons of proper theories of graphs must be a.e. $\{0, 1\}$ -valued (Theorem 6.2) and show that primally almost finite families of graphs must have bounded VC dimension (Theorem 6.10). In Section 7, we prove the general characterization of strongly persistent classes of structures in finite relational languages (Theorem 7.9). In the brief Section 8, we point out which results concerning WR generalize easily to finite relational languages. In the final Section 9, we summarize and discuss some open problems.

Acknowledgments

We thank two anonymous reviewers for helpful comments.

2 Preliminaries

In this section, we establish the notation and background results that will be used throughout the paper. The core results of the paper are in probabilistic combinatorics, and most of the results and proofs are stated in that language. Still, there are quite a few points where we believe the introduction of (simple) model theoretic language is more natural both to explain our approach and to organize the results, as we shall explain.

We denote the set of non-negative integers by \mathbb{N} and the set of positive integers by $\mathbb{N}_+ \stackrel{\text{def}}{=} \mathbb{N} \setminus \{0\}$ and given $n, k \in \mathbb{N}$, we let $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$ and let $(n)_k \stackrel{\text{def}}{=} n(n-1) \cdots (n-k+1)$ denote the falling factorial. Given a set V and $k \in \mathbb{N}$, we let $(V)_k$ be the set of *injective* functions $[k] \rightarrow V$, we let $\binom{V}{k} \stackrel{\text{def}}{=} \{A \subseteq V \mid |A| = k\}$ be the set of subsets of V of size k , let $\binom{V}{\leq k} \stackrel{\text{def}}{=} \bigcup_{\ell=0}^k \binom{V}{\ell}$ and we let $r(V) \stackrel{\text{def}}{=} \bigcup_{k \in \mathbb{N}_+} \binom{V}{k}$ be the set of non-empty finite subsets of V . We will often abuse notation and write n in place of $[n]$ when $V = [n]$ in some of the notation.

2.1 Terminology from model theory: structures and theories

In this paper a main object of study is hereditary classes of graphs. These can be seen as a special case of what are called in model theoretic language “classes of structures in finite relational languages” or even “universal theories”, and often the greater level of generality is useful. We now explain all these (quite natural) definitions.

Recall that a family of graphs (up to isomorphism) is called *hereditary* if it is closed under induced subgraphs. As an example, consider the triangle-free graphs, and observe the following. Using R as the binary edge symbol, we can write a set of axioms T_{Tri} of first order logic which capture this class of graphs. First, the theory of graphs T_{Graph} will say: the edge relation is symmetric $[\forall x \forall y (R(x, y) \iff R(y, x))]$ and irreflexive $[\forall x \forall y (R(x, y) \implies x \neq y)]$. To obtain T_{Tri} , we add the axiom that there are no triangles $[\forall x \forall y \forall z \neg(x \neq y \wedge y \neq z \wedge z \neq x \wedge R(x, y) \wedge R(y, z) \wedge R(x, z))]$. This T_{Tri} is called a universal theory because it uses only universal quantifiers, and to a model theorist, this explains the fact that the

axioms still hold on any induced substructure, or in other words, the graphs satisfying the axioms T_{Tri} form a hereditary class. Model theorists consider a set of axioms and the class of structures satisfying those axioms to be two sides of the same coin, so in logical parlance we could say we are studying the hereditary family of triangle-free graphs, or equivalently, we are studying the universal theory T_{Tri} .

To motivate the phrase “finite relational languages”, observe that there are other natural hereditary classes we might want to study, such as: linear orders, tournaments, 3-uniform hypergraphs; or perhaps the class of hypergraphs on which we have both a graph edge E and a 3-uniform hyperedge R , and E has no triangles and R has no tetrahedra (i.e., there are no four vertices such that every three form an R -hyperedge).¹ The following three definitions give us the right level of generality. First, we choose our alphabet.

Definition 2.1. A *finite relational language* \mathcal{L} is a set of finitely many symbols P_1, \dots, P_n , each given with an *arity* $k(P_i) \in \mathbb{N}_+$.

Second, when we define a graph G , we present it as a set V of vertices along with a set $R \subseteq V \times V$ of edges, and \mathcal{L} -structures just extend this in the obvious way:

Definition 2.2. Given a finite relational language \mathcal{L} , an \mathcal{L} -structure M is given by:

- (a) the data of a set $V(M)$, called the vertices of M or the domain of M , and
- (b) for each $P_i \in \mathcal{L}$, a subset of $V(M)^{k(P_i)}$, that is, the set of $k(P_i)$ -tuples on which P_i holds. This set is denoted P_i^M and called the interpretation of P_i in M .

Finally, we make the bridge to theories:²

Definition 2.3. A *universal theory* T in the language \mathcal{L} is a set of axioms (i.e., a set of well formed formulas of first order logic, using basic logical symbols along with the symbols from \mathcal{L}) in which the only quantifiers are universal. An \mathcal{L} -structure M is said to be a *model* for T , in symbols $M \models T$, if all the axioms T hold in M .³

Throughout this text, unless explicitly mentioned otherwise, all languages are assumed to be finite relational languages. We allow⁴ structures to have empty vertex sets and the unique structure with empty vertex set, called *trivial structure*, is denoted K_0 . Given an \mathcal{L} -structure M , $V \subseteq V(M)$ and $v \in V(M)$, we denote the *substructure* of M induced by V by $M|_V$ (i.e., we have $V(M|_V) \stackrel{\text{def}}{=} V$ and $P^{M|_V} \stackrel{\text{def}}{=} P^M \cap V^{k(P)}$ for every $P \in \mathcal{L}$) and we let $M - v \stackrel{\text{def}}{=} M|_{V(M) \setminus \{v\}}$.

We put the following in a convention environment to emphasize its importance:

Convention 2.4. Our substructures and subgraphs will always be induced, but keeping with the tradition of the fields, we will use the short term “substructure” for the former but the full term “induced subgraph” for the latter.

¹A series of formal examples will be worked out later in this section.

²Observe that each of the hereditary classes listed before Definition 2.1 can be expressed as a class of models for some appropriate universal axioms using an appropriate \mathcal{L} : for instance $\{\langle\cdot\rangle\}$ with $k(\langle\cdot\rangle) = 2$, $\{R\}$ with $k(R) = 2$, $\{R\}$ with $k(R) = 3$, and $\{E, R\}$ with $k(E) = 2$ and $k(R) = 3$, respectively.

³Formally defining “ $M \models T$ ” requires an induction on formula complexity, as in [CK90, Chapter 2].

⁴This may seem curious to model theorists but simplifies our calculations.

In this paper, and elsewhere, the word “structure” and “model” are used interchangeably; often the first emphasizes the abstract aspect, and the second emphasizes the relation to a theory. Basic model theory verifies that these work as intended: when T is a universal theory, the class of models of T is closed under substructures, i.e., is a hereditary class.

Convention 2.5. In this paper the phrase “universal theory” will denote both a fixed universal set of axioms T in some fixed finite relational language \mathcal{L} , and the *class of \mathcal{L} -structures* which are models for T , which, as noted, is a hereditary class.

The reader is free to substitute the phrase “hereditary class” for “universal theory” throughout, keeping in mind the language being used and the appropriate notion of substructure, and Convention 2.4.

Discussion 2.6. If the reader is essentially free to read “universal theory” as “hereditary class,” why do we introduce this terminology? This indicates a certain change in perspective which appears to be useful for theorems and proofs. Part of this choice reflects a history of work in the area, as in the theons of [CR23]. Centrally for the present work, a characteristically model-theoretic move of “studying all models of a theory” can be seen in the definition of the class \mathbf{WR} and in various aspects of the proofs.

2.2 Counting embeddings of graphs and structures

In order to develop the theory of graph limits (see 2.6 below), one starts by defining *labeled (induced) density of a finite graph H in some other finite graph G* . That is, let $T_{\text{ind}}(H, G)$ be the set of injective maps $f: V(H) \rightarrow V(G)$ which preserve edges and non-edges, and let

$$t_{\text{ind}}(H, G) \stackrel{\text{def}}{=} \begin{cases} \frac{|T_{\text{ind}}(H, G)|}{(|G|)^{|H|}}, & \text{if } |H| \leq |G|, \\ 0, & \text{otherwise.} \end{cases}$$

One also defines the *induced density of H in G* :

$$p(H, G) \stackrel{\text{def}}{=} \frac{|\{U \subseteq V(H) \mid G|_U \cong H\}|}{\binom{|H|}{|G|}} = \frac{|H|!}{|\text{Aut}(H)|} \cdot t_{\text{ind}}(H, G),$$

when $|H| \leq |G|$ (and defined as 0 otherwise), which gives the normalized number of (unlabeled) copies. The discussion in the previous subsection suggests an obvious extension of this definition to the setting of finite relational languages. That is, given finite structures M and N in a language \mathcal{L} , we let $T_{\text{ind}}(M, N)$ be the set of embeddings of M in N (i.e., the set of injective maps $f: V(M) \rightarrow V(N)$ that preserve all relations and their negations) and let

$$t_{\text{ind}}(M, N) \stackrel{\text{def}}{=} \begin{cases} \frac{|T_{\text{ind}}(M, N)|}{(|N|)^{|M|}}, & \text{if } |M| \leq |N|, \\ 0, & \text{otherwise} \end{cases}$$

be the *labeled (induced) density* of M in N . We also define the *(induced) density* of M in N as the normalized number of substructures of N that are isomorphic to M given by

$$p(M, N) \stackrel{\text{def}}{=} \frac{|\{U \subseteq V(N) \mid N|_U \cong M\}|}{\binom{|N|}{|M|}} = \frac{|M|!}{|\text{Aut}(M)|} \cdot t_{\text{ind}}(M, N),$$

when $|M| \leq |N|$ (and defined as 0 otherwise), where $\text{Aut}(M)$ is the group of automorphisms of M .

Given a universal theory T in \mathcal{L} and a set V , we let $\mathcal{K}_V[T]$ be the set of all models M of T whose vertex set $V(M)$ is V . Given $n \in \mathbb{N}$, we let $\mathcal{M}_n[T]$ be the set of models of T of size n up to isomorphism; we typically think of $\mathcal{M}_n[T]$ as a subset of $\mathcal{K}_{[n]}[T]$ by putting one representative of each isomorphism class in $\mathcal{M}_n[T]$. We also let $\mathcal{M}[T] \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} \mathcal{M}_n[T]$.

2.3 Open interpretations

When comparing hereditary classes it is useful to know when one contains the information of the other in perhaps a different presentation. As a trivial example, consider the hereditary class of triangle-free graphs in which the edge relation is called E ($k(E) = 2$) and the hereditary class of triangle-free graphs in which the edge relation is called R ($k(R) = 2$). As a slightly less trivial example, compare these to the hereditary class of graphs in which each vertex is either colored red or green, and the green vertices form a triangle-free graph.

This subsection introduces language for addressing such situations by identifying obviously equivalent pieces of hereditary classes. Model theorists will recognize “open” as meaning “quantifier-free”.

Recall that for universal theories T_1 and T_2 in finite relational languages \mathcal{L}_1 and \mathcal{L}_2 , respectively, an *open interpretation* (or *definition*) from T_1 to T_2 is a function I (denoted $I: T_1 \rightsquigarrow T_2$) that maps each predicate symbol $P \in \mathcal{L}_1$ to an open (i.e., quantifier-free) formula $I(P)(x_1, \dots, x_{k(P)})$ in \mathcal{L}_2 and such that for each axiom $\forall \vec{x}, F(\vec{x})$ of T_1 , we have $T_2 \vdash \forall \vec{x}, I(F)(\vec{x})$ when we declare I to commute with logical connectives. Open interpretations of the form $I: T_1 \rightsquigarrow T_2$ contra-variantly define maps $\mathcal{K}_V[T_2] \rightarrow \mathcal{K}_V[T_1]$ for each set V given by $(I(M) \models P(\vec{x})) \iff (M \models I(P)(\vec{x}))$ for each $P \in \mathcal{L}_1$. In turn, for an open interpretation $I: T_1 \rightsquigarrow T_2$, we let $I(T_2)$ be the universal theory in the language of T_1 whose finite models are precisely those of the form $I(M)$ for some $M \in \mathcal{M}[T_2]$, that is, the axioms of $I(T_2)$ are

$$\forall x_1, \dots, x_n, \bigvee_{M \in \mathcal{K}_n[T_2]} D_{\text{open}}(I(M))(x_1, \dots, x_n) \quad (n \in \mathbb{N}),$$

where $D_{\text{open}}(N)$ is the *open diagram* of N , that is, the open formula

$$\bigwedge_{1 \leq i < j \leq n} x_i \neq x_j \wedge \bigwedge_{P \in \mathcal{L}_2} \left(\bigwedge_{\alpha \in P^N} P(x_{\alpha_1}, \dots, x_{\alpha_{k(P)}}) \wedge \bigwedge_{\alpha \in V(N)^{k(P)} \setminus P^N} \neg P(x_{\alpha_1}, \dots, x_{\alpha_{k(P)}}) \right)$$

that completely encodes the quantifier-free type (over \emptyset) of the tuple $(1, \dots, n)$ in N . To make sense out of $D_{\text{open}}(K_0)$ (which must be a quantifier-free formula on zero variables), we allow our formulas to use the *tautological truth symbol* \top so that $D_{\text{open}}(K_0)$ is defined as \top .

An *INT-isomorphism* (or *interdefinition*) is an open interpretation $I: T_1 \rightsquigarrow T_2$ such that there exists an open interpretation $J: T_2 \rightsquigarrow T_1$ such that for every set V , the compositions $J \circ I: \mathcal{K}_V[T_2] \rightarrow \mathcal{K}_V[T_2]$ and $I \circ J: \mathcal{K}_V[T_1] \rightarrow \mathcal{K}_V[T_1]$ are the identity maps. Since $p(M, N) = p(I(M), I(N))$ whenever $I: T_1 \rightsquigarrow T_2$ is an INT-isomorphism and M and N are finite models of T_2 , we typically do not distinguish between INT-isomorphic universal theories.

2.4 Canonical theories: avoiding the diagonal

When defining graphs we require the edge relation to be irreflexive. In general, a universal theory T in \mathcal{L} is *canonical* if T entails

$$\forall x_1, \dots, x_{k(P)}, \left(\bigvee_{i \neq j} x_i = x_j \rightarrow \neg P(\vec{x}) \right) \quad (1)$$

for every predicate symbol $P \in \mathcal{L}$. By [CR20, Theorem 2.3] (see also [AC14, §2.2]), every universal theory is INT-isomorphic to some canonical theory and as such, from this point forward, all theories are assumed to be canonical theories, unless explicitly mentioned otherwise.

We say that a canonical theory T is *non-degenerate* if it contains some infinite model (equivalently, if $\mathcal{M}_n[T]$ is non-empty for every $n \in \mathbb{N}$).

2.5 Examples of theories used in the text

To make more concrete our use of “universal theories” rather than simply hereditary classes of graphs, in this subsection we lay out some of the main examples used in the text, along with a useful construction of a canonical theory. We have already seen the *theory of graphs* T_{Graph} , and we will say that T “is a universal theory of graphs” when T is a universal theory that is obtained from T_{Graph} by adding axioms, that is, its finite models are some hereditary class of graphs (an obvious example is T_{Graph} itself, a less obvious one is the theory T_{Tri} of triangle-free graphs).

A second kind of example is the *theory of k -hypergraphs* $T_{k\text{-Hypergraph}}$, that is, the canonical theory with a single *symmetric irreflexive*⁵ k -ary predicate E . Obviously the theory of graphs is simply $T_{\text{Graph}} = T_{2\text{-Hypergraph}}$. In these theories, we denote by $K_n^{(k)} \in \mathcal{M}_n[T_{k\text{-Hypergraph}}]$ the *complete k -hypergraph on n vertices* (i.e., we have $V(K_n^{(k)}) \stackrel{\text{def}}{=} [n]$ and $E^{K_n^{(k)}} \stackrel{\text{def}}{=} ([n])_k$) and we let $K_n \stackrel{\text{def}}{=} K_n^{(2)}$ be the *complete graph on n vertices*. Given a k -hypergraph G , we let \overline{G} denote the *complement* hypergraph of G (given by $V(\overline{G}) \stackrel{\text{def}}{=} V(G)$ and $E^{\overline{G}} \stackrel{\text{def}}{=} ([n])_k \setminus E^G$). In particular, $\overline{K}_n^{(k)}$ is the *empty k -hypergraph on n vertices*.

⁵In the sense that the predicate is *not* true in any non-injective tuple.

Another example of a canonical theory is the *theory of (strict) linear orders* T_{LinOrder} , i.e., the theory with a binary predicate symbol \prec and axioms

$$\begin{aligned} & \forall x, \neg(x \prec x), \\ & \forall \vec{x}, (x_1 \neq x_2 \rightarrow (x_1 \prec x_2 \vee x_2 \prec x_1)), \\ & \forall \vec{x}, (x_1 \prec x_2 \wedge x_2 \prec x_3 \rightarrow x_1 \prec x_3). \end{aligned}$$

Another useful theory is the *theory of permutations* $T_{\text{Perm}} \stackrel{\text{def}}{=} T_{\text{LinOrder}} \cup T_{\text{LinOrder}}$, which is the theory of two (strict) linear orders on the same ground set. Its finite models (up to isomorphism) are in one-to-one correspondence to usual permutations via $S_n \ni \sigma \mapsto M_\sigma \in \mathcal{M}_n[T_{\text{Perm}}]$, where the first order $\prec_1^{M_\sigma}$ of M_σ is simply the natural order on $[n]$ and the second order is given by

$$(M_\sigma \models i \prec_2 j) \iff \sigma^{-1}(i) < \sigma^{-1}(j).$$

Some other examples of theories that can be obtained from T_{Graph} by adding axioms that will be used are the *theory of graphs of agreements of permutations* $T_{\text{PermGraph}} \stackrel{\text{def}}{=} I(T_{\text{Perm}})$, where $I: T_{\text{Graph}} \rightsquigarrow T_{\text{Perm}}$ is given by

$$I(E)(x, y) \stackrel{\text{def}}{=} (x \neq y \wedge (x \prec_1 y \leftrightarrow x \prec_2 y)),$$

the *theory of bipartite graphs* $T_{\text{Bipartite}}$, which is obtained from T_{Graph} by adding the axioms $\forall \vec{x}, \neg D_{\text{open}}(C_{2n+1})(\vec{x})$ for every $n \in \mathbb{N}_+$, where C_ℓ is the ℓ -cycle graph and the *theory of perfect graphs*, which by the Strong Perfect Graph Theorem [CRST06], is obtained from T_{Graph} by adding the axioms $\forall \vec{x}, \neg(D_{\text{open}}(C_{2n+1})(\vec{x}) \vee D_{\text{open}}(\bar{C}_{2n+1})(\vec{x}))$ for every $n \geq 2$.

For every finite relational language \mathcal{L} , we let $T_{\mathcal{L}}$ be the *pure canonical theory* in \mathcal{L} , that is, the theory whose axioms are precisely (1) for each $P \in \mathcal{L}$. Unless explicitly mentioned otherwise, all \mathcal{L} -structures are assumed to be *canonical structures*, that is, models of $T_{\mathcal{L}}$.

Given a family \mathcal{F} of models of a canonical theory T , we let $\text{Forb}_T(\mathcal{F})$ be the theory of models of T that do not have any copies of models in \mathcal{F} , that is, $\text{Forb}_T(\mathcal{F})$ is obtained from T by adding the axioms $\forall \vec{x}, \neg D_{\text{open}}(F)(\vec{x})$ for every $F \in \mathcal{F}$ (note that if $F = K_0$, then this formula takes the form $\neg \top$, which is tautologically false, so $\text{Forb}_T(\mathcal{F})$ has no models when $K_0 \in \mathcal{F}$).

Another simple but useful construction is the following: if \mathcal{F} is a family of finite \mathcal{L} -structures that is closed under substructures, then we let $\text{Th}(\mathcal{F}) \stackrel{\text{def}}{=} \text{Forb}_{T_{\mathcal{L}}}(\mathcal{M}[T_{\mathcal{L}}] \setminus \mathcal{F})$ be the unique universal theory (up to reaxiomatization) such that $\mathcal{M}[\text{Th}(\mathcal{F})] = \mathcal{F}$.

2.6 Basics of graphons and theons

We start here with the language of models before specializing to the case of graphons, which will be central for an initial segment of the text.

A sequence $(N_n)_{n \in \mathbb{N}}$ of finite models of a canonical theory T is called *convergent* if it is *increasing* in the sense that $|N_n| < |N_{n+1}|$ for every $n \in \mathbb{N}$ and if for every $M \in \mathcal{M}[T]$,

the limit $\lim_{n \rightarrow \infty} p(M, N_n)$ exists. Another way of seeing this convergence is that each finite model N of T corresponds to the point $p(-, N) \in [0, 1]^{\mathcal{M}[T]}$ and convergence of an increasing sequence $(N_n)_{n \in \mathbb{N}}$ amounts to convergence in the (compact metrizable) product topology of $[0, 1]^{\mathcal{M}[T]}$ of the corresponding sequence $(p(-, N_n))_{n \in \mathbb{N}}$.

There are essentially two ways to encode limits of convergent sequences. The first is algebraically/syntactically: we say that a function $\phi: \mathcal{M}[T] \rightarrow [0, 1]$ is the limit of a convergent sequence $(N_n)_{n \in \mathbb{N}}$ if $\phi(M) = \lim_{n \rightarrow \infty} p(M, N_n)$ for every $M \in \mathcal{M}[T]$. The theory of flag algebras [Raz07] describes the set $\text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ of functions that are limits of convergent sequences as precisely as the ones that induce positive homomorphisms from a particular commutative \mathbb{R} -algebra $\mathcal{A}[T]$ to \mathbb{R} , but for this work, the unfamiliarized reader can safely think of $\text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ as simply a fancy notation for the subset of $[0, 1]^{\mathcal{M}[T]}$ of all ϕ that are limits of some convergent sequence. Note that compactness of $[0, 1]^{\mathcal{M}[T]}$ implies that $\text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is non-empty if and only if T is non-degenerate.

For $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$, the *theory of positive models* of ϕ is the universal theory $\text{Th}(\phi)$ whose finite models are precisely those models M of T such that $\phi(M) > 0$, that is, the axioms of $\text{Th}(\phi)$ are

$$\forall \vec{x}, \bigvee_{\substack{M \in \mathcal{K}_n[T] \\ \phi(M) > 0}} D_{\text{open}}(M)(x_1, \dots, x_n) \quad (n \in \mathbb{N}).$$

The second way of encoding limits is geometrically/semantically. In the case of graphs, we can encode limits using a *graphon* W over an atomless standard probability space $\Omega = (X, \mathcal{A}, \mu)$, that is, W is a symmetric function $X \times X \rightarrow [0, 1]$ measurable in the completion of the product σ -algebra (typically, the space Ω is taken to be $[0, 1]$ equipped with the Lebesgue measure λ , in which case, a graphon is simply a symmetric Lebesgue measurable function $[0, 1]^2 \rightarrow [0, 1]$). Given one such graphon W over $\Omega = (X, \mathcal{A}, \mu)$ and a finite graph G , the *labeled (induced) density* and the *(induced) density* of G in W are defined respectively as

$$t_{\text{ind}}(G, W) \stackrel{\text{def}}{=} \int_{X^{V(G)}} \prod_{\{v, w\} \in E(G)} W(x_v, x_w) \prod_{\{v, w\} \in E(\bar{G})} (1 - W(x_v, x_w)) d\mu(x),$$

$$\phi_W(G) \stackrel{\text{def}}{=} p(G, W) \stackrel{\text{def}}{=} \frac{|G|!}{|\text{Aut}(G)|} \cdot t_{\text{ind}}(G, W),$$

where $E(G) \stackrel{\text{def}}{=} \{\{v, w\} \mid G \models E(v, w)\}$ is the edge set of G and \bar{G} is the complement of G . We say that a convergent sequence $(H_n)_{n \in \mathbb{N}}$ of graphs converges to W if $\lim_{n \rightarrow \infty} p(G, H_n) = \phi_W(G)$ for every $G \in \mathcal{M}[T_{\text{Graph}}]$. Another way of interpreting $t_{\text{ind}}(G, W)$ above is to define the set $T_{\text{ind}}(G, W)$ of *labeled (induced) copies* of G in W as

$$T_{\text{ind}}(G, W) \stackrel{\text{def}}{=} \left\{ (x, y) \in X^{V(G)} \times [0, 1]^{\binom{V(G)}{2}} \mid \begin{array}{l} \forall \{v, w\} \in \binom{V(G)}{2}, \\ (\{v, w\} \in E(G) \leftrightarrow y_{\{v, w\}} < W(x_v, x_w)) \end{array} \right\} \quad (2)$$

and note that $t_{\text{ind}}(G, W) = (\mu \otimes \lambda)(T_{\text{ind}}(G, W))$. We also use the shorthand $\text{Th}(W) \stackrel{\text{def}}{=} \text{Th}(\phi_W)$ for the theory of positive graphs of W . Note that when W is $\{0, 1\}$ -valued, we can interpret it as the adjacency matrix of a graph with vertex set X and for $(x, y) \in X^{V(G)} \times [0, 1]^{\binom{V(G)}{2}}$ such that all coordinates of x are distinct, we have $(x, y) \in T_{\text{ind}}(G, W)$ if and only if x is an embedding of G in the graph encoded by the $\{0, 1\}$ -valued W .

The main theorem of the theory of graphons [LS06] says that graphons precisely encode limits of convergent graph sequences. Along with the flag algebra description, this can be easily summarized as $\text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R}) = \{\phi_W \mid W \text{ is a graphon}\}$. However, let us note that different graphons can represent the same limit; for example, for any graphon W over $[0, 1]$, the graphon W' defined by $W'(x, y) \stackrel{\text{def}}{=} W(2x \bmod 1, 2y \bmod 1)$ represents the same limit as W (i.e., we have $\phi_W = \phi_{W'}$).

Another very useful theorem is the Graphon Removal Lemma [Pet13, Theorem 1], which says that for any graphon W over Ω , there exists a graphon W' that differs from W only by a zero measure set (hence $\phi_W = \phi_{W'}$) and such that for every $G \in \mathcal{M}[T_{\text{Graph}}]$, if $t_{\text{ind}}(G, W) = 0$, then $T_{\text{ind}}(G, W') \subseteq \mathcal{D}_{V(G)}$, where

$$\mathcal{D}_V \stackrel{\text{def}}{=} \{(x, y) \in X^V \times [0, 1]^{\binom{V}{2}} \mid \exists v, w \in V, (v \neq w \wedge x_v = x_w)\} \quad (3)$$

is the *diagonal set*, that is, the Graphon Removal Lemma says that we only need to change W in a zero measure set to remove all off-diagonal copies of finite graphs that have zero density in W .

For the general case, we will use the theory of theons [CR20] (see also [Aus08] and [AC14] for alternative semantic limits). Given an atomless standard probability space $\Omega = (X, \mathcal{A}, \mu)$ and a set V , we let $\mathcal{E}_V(\Omega) \stackrel{\text{def}}{=} X^{r(V)}$ (recall that $r(V) \stackrel{\text{def}}{=} \bigcup_{k \in \mathbb{N}_+} \binom{V}{k}$), equipping it with the completion of the product measure, which is denoted μ as well, by abuse.

For a predicate symbol P in a finite relational language \mathcal{L} , a P -on over Ω is a measurable subset of $\mathcal{E}_k(P)(\Omega)$. An *Euclidean structure* in \mathcal{L} over Ω is a function \mathcal{N} that maps each predicate symbol $P \in \mathcal{L}$ to a P -on $\mathcal{N}_P \subseteq \mathcal{E}_k(P)(\Omega)$. If we are further given a finite (canonical) \mathcal{L} -structure K , we define the set of *labeled (induced) copies* of K in \mathcal{N} as

$$T_{\text{ind}}(K, \mathcal{N}) \stackrel{\text{def}}{=} \bigcap_{P \in \mathcal{L}} \left(\bigcap_{\alpha \in P^K} (\alpha^*)^{-1}(\mathcal{N}_P) \cap \bigcap_{\alpha \in (V(K))_{k(P)} \setminus P^K} (\alpha^*)^{-1}(\mathcal{E}_{k(P)}(\Omega) \setminus \mathcal{N}_P) \right),$$

where for each injection $\alpha: [k] \rightarrow V$, $\alpha^*: \mathcal{E}_V(\Omega) \rightarrow \mathcal{E}_k(\Omega)$ is the contra-variantly defined “projection” given by

$$\alpha^*(x)_A \stackrel{\text{def}}{=} x_{\alpha(A)} \quad (x \in \mathcal{E}_V(\Omega), A \in r(k)). \quad (4)$$

Similarly to the graphon case, we let

$$t_{\text{ind}}(K, \mathcal{N}) \stackrel{\text{def}}{=} \mu(T_{\text{ind}}(K, \mathcal{N})), \quad \phi_{\mathcal{N}}(K) \stackrel{\text{def}}{=} \frac{|K|!}{|\text{Aut}(K)|} \cdot t_{\text{ind}}(K, \mathcal{N}),$$

and we say that a convergent sequence of finite structures $(N_n)_{n \in \mathbb{N}}$ converges to \mathcal{N} if $\lim_{n \rightarrow \infty} p(K, N_n) = \phi_{\mathcal{N}}(K)$ for every finite structure K . Similarly, we use the shorthand $\text{Th}(\mathcal{N}) \stackrel{\text{def}}{=} \text{Th}(\phi_{\mathcal{N}})$ for the theory of positive models of \mathcal{N} .

For a canonical theory T in \mathcal{L} , a *(weak) T -on* over Ω is an Euclidean structure \mathcal{N} in \mathcal{L} over Ω such that $\phi_{\mathcal{N}}(K) = 0$ whenever K is a finite \mathcal{L} -structure that is *not* a model of T . A *strong T -on* over Ω is a T -on \mathcal{N} such that for every finite \mathcal{L} -structure K that is *not* a model of T , we have $T_{\text{ind}}(K, \mathcal{N}) \subseteq \mathcal{D}_{V(K)}(\Omega)$, where

$$\mathcal{D}_V(\Omega) \stackrel{\text{def}}{=} \{x \in \mathcal{E}_V(\Omega) \mid \exists v, w \in V, (v \neq w \wedge x_{\{v\}} = x_{\{w\}})\} \quad (5)$$

denotes the *diagonal set*.

The main theorem of the theory of theons says that theons precisely encode limits of convergent sequences of models, that is, we have $\text{Hom}^+(\mathcal{A}[T], \mathbb{R}) = \{\phi_{\mathcal{N}} \mid \mathcal{N} \text{ is a } T\text{-on}\}$. In fact, the easy inclusion of this equality is worth spelling out: given a T -on \mathcal{N} over Ω , for each $n \in \mathbb{N}$, we sample $\boldsymbol{\theta}$ in $\mathcal{E}_n(\Omega)$ according to μ and let \mathbf{N}_n be the random element of $\mathcal{K}_n[\mathcal{T}_{\mathcal{L}}]$ given by

$$\begin{aligned} V(\mathbf{N}_n) &\stackrel{\text{def}}{=} [n], \\ (\mathbf{N}_n \vDash P(\alpha)) &\iff \alpha^*(\boldsymbol{\theta}) \in \mathcal{N}_P \quad (P \in \mathcal{L}, \alpha \in ([n])_{k(P)}), \end{aligned}$$

where $\alpha^*: [k(P)] \rightarrow [n]$ is given by (4). It is a straightforward exercise on distribution concentration to check that with probability 1, the sequence $(\mathbf{N}_n)_{n \in \mathbb{N}}$ converges to $\phi_{\mathcal{N}}$. In particular, this means that any limit $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is also a limit of a sequence of models $(N_n)_{n \in \mathbb{N}}$ that does not omit sizes in the sense that $|N_n| = n$ for every $n \in \mathbb{N}$. However, similarly to graphons, different theons can represent the same limit.

Similarly to graphons, another very useful theorem of the theory of theons is the Induced Euclidean Removal Lemma [CR20, Theorem 3.3], which says that any weak T -on can be turned into a strong T -on by changing its peons only on a zero measure set (which in particular means the two theons represent the same limit). A fortiori, by viewing a T -on \mathcal{N} as a $\text{Th}(\mathcal{N})$ -on, the Induced Euclidean Removal Lemma implies that there exists a T -on \mathcal{N}' whose peons differ from those of \mathcal{N} only by a zero measure set and such that $T_{\text{ind}}(K, \mathcal{N}') \subseteq \mathcal{D}_{V(K)}(\Omega)$ whenever $t_{\text{ind}}(K, \mathcal{N}') = 0$.

Given an open formula $F(x_1, \dots, x_n)$ in \mathcal{L} and an Euclidean structure \mathcal{N} in \mathcal{L} over Ω , the *truth set* $T(F, \mathcal{N}) \subseteq \mathcal{E}_n(\Omega)$ of F in \mathcal{N} is defined inductively as follows.

- i. $T(x_i = x_i, \mathcal{N}) \stackrel{\text{def}}{=} \mathcal{E}_n(\Omega).$
- ii. $T(x_i = x_j, \mathcal{N}) \stackrel{\text{def}}{=} \emptyset$, if $i \neq j$.
- iii. $T(P(x_{\alpha_1}, \dots, x_{\alpha_{k(P)}}), \mathcal{N}) \stackrel{\text{def}}{=} \emptyset$, if $\alpha: [k(P)] \rightarrow [n]$ is not injective.
- iv. $T(P(x_{\alpha_1}, \dots, x_{\alpha_{k(P)}}), \mathcal{N}) \stackrel{\text{def}}{=} (\alpha^*)^{-1}(\mathcal{N}_P)$ if $\alpha: [k(P)] \rightarrow [n]$ is injective, where α^* is as in (4) for $V = [n]$.

v. $T(-, \mathcal{N})$ commutes with logical connectives (so e.g., $T(\neg F, \mathcal{N}) \stackrel{\text{def}}{=} \mathcal{E}_n(\Omega) \setminus T(F, \mathcal{N})$ and $T(F_1 \vee F_2, \mathcal{N}) \stackrel{\text{def}}{=} T(F_1, \mathcal{N}) \cup T(F_2, \mathcal{N})$).

One might argue that items (ii) and (iii) above should be defined as particular subsets of the diagonal $\mathcal{D}_n(\Omega)$ (see (5)), but since all information is lost in $\mathcal{D}_n(\Omega)$ (both in weak and strong theons), the definition above is just as good but simpler. It is straightforward to check that for $K \in \mathcal{K}_n[T_{\mathcal{L}}]$, we have $T_{\text{ind}}(K, \mathcal{N}) = T(D_{\text{open}}(K), \mathcal{N})$.

Open interpretations behave very well with the notion of convergence. Furthermore, there are natural operations in the theories of flag algebras [Raz07, Definition 4 and Theorem 2.6] and theons [CR20, Remark 6] that capture this action in the limit. Namely, if $I: T_1 \rightsquigarrow T_2$ is an open interpretation and $(N_n)_{n \in \mathbb{N}}$ is a convergent sequence of finite models of T_2 converging to $\phi \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ and to the T_2 -on \mathcal{N} , then $(I(N_n))_{n \in \mathbb{N}}$ is a convergent sequence of models of T_1 converging to $\phi^I \in \text{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and to the T_1 -on $I(\mathcal{N})$ given by

$$\begin{aligned} \phi^I(M) &\stackrel{\text{def}}{=} \sum_{\substack{M' \in \mathcal{M}[T_2] \\ I(M') \cong M}} \phi(M') \quad (M \in \mathcal{M}[T_1]), \\ I(\mathcal{N})_P &\stackrel{\text{def}}{=} T(I(P), \mathcal{N}) \quad (P \in \mathcal{L}). \end{aligned} \quad (6)$$

Given $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$, a *sub-object* of ϕ of measure $c > 0$ is a $\psi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ such that there exist a sequence $(N_n)_{n \in \mathbb{N}}$ of models converging to ϕ and sets $A_n \subseteq V(N_n)$ such that $\lim_{n \rightarrow \infty} |A_n|/|N_n| = c$ and $(N_n|_{A_n})_{n \in \mathbb{N}}$ converges to ψ . By a small abuse, we may use theons \mathcal{N} and \mathcal{H} in place of ϕ and/or ψ , respectively when $\phi_{\mathcal{N}} = \phi$ and $\phi_{\mathcal{H}} = \psi$. When the underlying theory is T_{Graph} , we will use the more natural name *subgraphon* for the concept of sub-object and with a similar abuse, we will use graphons W and W' in place of ϕ and/or ψ , respectively when $\phi_W = \phi$ and $\phi_{W'} = \psi$.

For simplicity and for later quotation we spell this out:

Definition 2.7 (Subgraphons). Given a graphon W over an atomless standard probability space Ω , a (positive measure) subgraphon W' of W is a graphon over a space Ω' such that there exist a sequence $(H_n)_{n \in \mathbb{N}}$ of graphs converging to W and sets $U_n \subseteq V(H_n)$ such that $\lim_{n \rightarrow \infty} |U_n|/|H_n| > 0$ and $(H_n|_{U_n})_{n \in \mathbb{N}}$ converges to W' .

Some useful equivalences are the following.

By [CM22, Lemma 3.3], $\psi \in \text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$ is a subgraphon of W of measure $c > 0$ if and only if there exists a measurable function $f: X \rightarrow [0, 1]$ with $\int_X f \, d\mu = c$ such that $\psi = \phi_{W|_f}$, where $W|_f$ is the graphon over the space $\Omega_f \stackrel{\text{def}}{=} (X, \mathcal{A}, \mu_f)$ defined by

$$\begin{aligned} \mu_f(B) &\stackrel{\text{def}}{=} \frac{1}{c} \int_B f(x) \, d\mu(x), \\ W|_f(x, y) &\stackrel{\text{def}}{=} W(x, y). \end{aligned} \quad (7)$$

More generally, by [CM22, Lemma 5.8], $\psi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is a sub-object of a theon \mathcal{N} over $\Omega = (X, \mathcal{A}, \mu)$ of measure $c > 0$ if and only if there exist a measurable function

$f: X \rightarrow [0, 1]$ with $\int_X f \, d\mu = c$ and a measure-isomorphism F modulo 0 from the space $\Omega_f = (X, \mathcal{A}, \mu_f)$ given by (7) to Ω such that $\psi = \phi_{\mathcal{N}|_f^F}$ for the theon $\mathcal{N}|_f^F$ over Ω_f defined by

$$(\mathcal{N}|_f^F)_P \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{k(P)}(\Omega_f) \mid x^F \in \mathcal{N}_P\} \quad (P \in \mathcal{L}),$$

where $x^F \in \mathcal{E}_{k(P)}(\Omega)$ is given by

$$x_B^F \stackrel{\text{def}}{=} \begin{cases} x_B, & \text{if } |B| = 1, \\ F(x_B), & \text{if } |B| \geq 2. \end{cases}$$

When the function f in the above is the indicator function $\mathbb{1}_A$ of some positive measure set $A \subseteq X$, we use the shorthands $\mu_A \stackrel{\text{def}}{=} \mu_{\mathbb{1}_A}$, $\Omega_A \stackrel{\text{def}}{=} \Omega_{\mathbb{1}_A}$, $W|_A \stackrel{\text{def}}{=} W|_{\mathbb{1}_A}$ and $\mathcal{N}|_A^F \stackrel{\text{def}}{=} \mathcal{N}|_{\mathbb{1}_A}^F$ for the concepts above. However, we point out that not every sub-object of \mathcal{N} is necessarily of the form $\mathcal{N}|_A^F$ for some positive measure set A (see [CR23, Example 45]).

2.7 The approximate Erdős–Hajnal property

We conclude this section by recalling the definition of the approximate Erdős–Hajnal property (**AEHP**) from [CM22, Definition 7.1].

Definition 2.8. A universal theory T in a finite relational language \mathcal{L} has the approximate Erdős–Hajnal property (**AEHP**) if every limit ϕ of T has a *trivial* sub-object, i.e., a sub-object ψ of the form $\psi = \phi_{\mathcal{N}}$ for some T -on \mathcal{N} whose peons all have measure in $\{0, 1\}$.

In particular, the definition above specializes to universal theories of graphs as follows.

Definition 2.9. If T is a universal theory of graphs (in other words, if T is a hereditary class of graphs), then $T \in \text{AEHP}$ if every graphon that is a limit of T has a *trivial* subgraphon, i.e., a subgraphon that is either a.e. equal to 0 or a.e. equal to 1.

An equivalent formulation of **AEHP** (see [CM22, Theorem 7.11]) is that for every convergent sequence $(N_n)_{n \in \mathbb{N}}$ of models of T , there exist sets $U_n \subseteq V(N_n)$ with $\lim_{n \rightarrow \infty} |U_n|/|N_n| > 0$ and $(N_n|_{U_n})_{n \in \mathbb{N}}$ converges to a trivial limit. In other words, **AEHP** for graphs requires linear-sized almost cliques or almost anti-cliques in the presence of convergence.

The property **AEHP** was introduced and its graph version was characterized both combinatorially and model-theoretically in [CM22].

3 Substitution and primality

A simple operation for graphs that is useful in studying the Erdős–Hajnal property and its approximate version (**AEHP**, see Definitions 2.8 and 2.9) is substitution. While for graphs this operation has received considerable attention [Gia97, APS01, Zve03, CKOS16], we will need a slight generalization of it for structures in finite relational languages along with some associated notions (e.g., primality).

In the graph case, some of the results of this section have appeared in some shape or another in the literature. As such, in Section 3.1, we state the definitions and the results for graphs that we will use without proofs but with pointers to their corresponding generalizations for relational structures that appear in Section 3.2 with proofs. The reader that feels sufficiently confident in their knowledge of these and is only interested in the graph case may read Section 3.1 then freely skip the remainder of the section.

3.1 Substitution and primality for graphs

Definition 3.1 (Graph version of Definition 3.8). Given two graphs F_1 and F_2 and $v \in V(F_1)$, the *substitution* of v in F_1 by F_2 is the graph $F_1^{v \rightarrow F_2}$ obtained from the disjoint union of $F_1 - v$ with F_2 by adding all edges of the form $\{u, w\}$, where $u \in V(F_1)$ is a vertex that is adjacent to v in F_1 and w is a vertex of F_2 (see Figure 1 for an example).

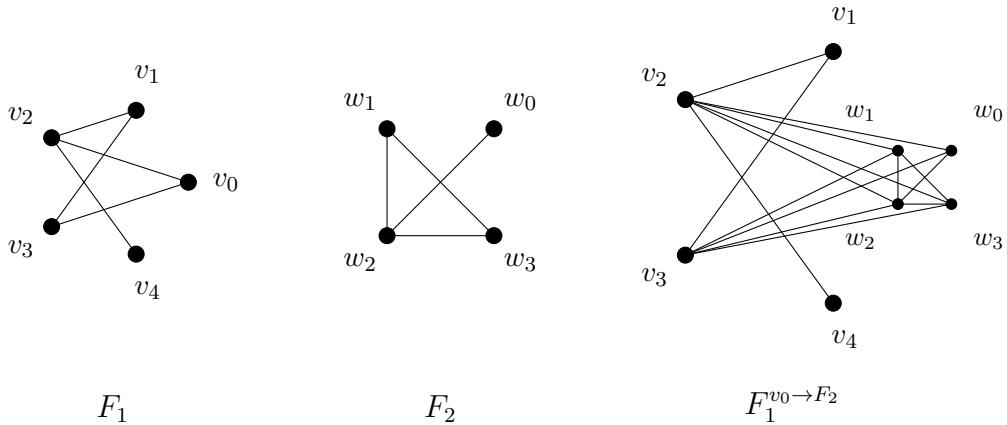


Figure 1: Example of a graph substitution.

We say a family \mathcal{F} of graphs (up to isomorphism) is *closed under substitution* if for every $F_1, F_2 \in \mathcal{F}$ and every $v \in V(F_1)$, we have $F_1^{v \rightarrow F_2} \in \mathcal{F}$. The *closure under substitutions* of \mathcal{F} is the smallest family $S(\mathcal{F})$ containing \mathcal{F} that is closed under substitution.

A graph F is called *prime* if it is *not* a substitution of v in F_1 by F_2 for any graphs F_1, F_2 and $v \in V(F_1)$ with $|F_1|, |F_2| < |F|$.

Lemma 3.2 (Graph version of Lemma 3.15). *Let \mathcal{F} be a non-empty family of graphs (up to isomorphism) that is closed under substitutions. Then \mathcal{F} is closed under induced subgraphs if and only if \mathcal{F} contains the trivial graph K_0 of size 0.*

Lemma 3.3 (Graph version of Lemma 3.16). *Let F_1, F_2 be finite graphs and let $v \in V(F_1)$.*

If F is an induced subgraph of $F_1^{v \rightarrow F_2}$, then there exist induced subgraphs G_1 and G_2 of F_1 and F_2 , respectively, with $v \in V(G_1)$ such that $F \cong G_1^{v \rightarrow G_2}$.

Conversely, if G_1 and G_2 are induced subgraphs of F_1 and F_2 , respectively, with $v \in V(G_1)$, then $G_1^{v \rightarrow G_2}$ is an induced subgraph of $F_1^{v \rightarrow F_2}$.

Lemma 3.4 (Graph version of Lemma 3.17). *If \mathcal{F} is a (not necessarily finite) family of graphs (up to isomorphism), $F \in S(\mathcal{F})$ and P is a prime induced subgraph of F , then P is an induced subgraph of some $F' \in \mathcal{F}$.*

Lemma 3.5 (Graph version of Lemma 3.18). *Let \mathcal{F} be a family of finite graphs (up to isomorphism) that is closed under substitutions and closed under induced subgraphs and let \mathcal{P} be the set of graphs in \mathcal{F} that are prime. Then $\mathcal{F} = S(\mathcal{P})$.*

Conversely, if \mathcal{P}' is a family of prime finite graphs that is closed under prime induced subgraphs and $\mathcal{F} = S(\mathcal{P}')$, then $\mathcal{P}' = \mathcal{P}$.

Definition 3.6 (Graph version of Definition 3.20). We say that a family of graphs \mathcal{F} (up to isomorphism) is *almost finite* if \mathcal{F} does not contain any infinite antichain in the induced subgraph partial order. Equivalently, \mathcal{F} is almost finite if for every infinite $\mathcal{F}' \subseteq \mathcal{F}$, there exist $F_1, F_2 \in \mathcal{F}'$ such that F_1 is a proper induced subgraph of F_2 .

By letting further \mathcal{P} be the set of all graphs of \mathcal{F} that are prime, we say that \mathcal{F} is *primally finite* if \mathcal{P} is finite and we say that \mathcal{F} is *primally almost finite* if \mathcal{P} is almost finite.

Lemma 3.7 (Graph version of Lemma 3.23). *The following are equivalent for a family \mathcal{F} of finite graphs (up to isomorphism).*

- i. *The family \mathcal{F} is almost finite.*
- ii. *For every sequence $(F_n)_{n \in \mathbb{N}}$ in \mathcal{F} , there exist $n, m \in \mathbb{N}$ such that $n < m$ and F_n is an induced subgraph of F_m .*

3.2 Substitution and primality for relational structures

Definition 3.8. Given two structures F_1 and F_2 in a finite relational language \mathcal{L} and $v \in V(F_1)$, a *substitution* of v in F_1 by F_2 is an \mathcal{L} -structure F such that there exist functions $f_1: V(F_1 - v) \rightarrow V(F)$ and $f_2: V(F_2) \rightarrow V(F)$ such that

- i. $V(F) = \text{im}(f_1) \cup \text{im}(f_2)$,
- ii. f_2 is an embedding of F_2 in F ,
- iii. For every $u \in V(F_2)$, the extension of f_1 to a function $V(F_1) \rightarrow V(F)$ that maps v to $f_2(u)$ is an embedding of F_1 in F .

We call the substitution F *standard* if $V(F_1) \cap V(F_2) = \emptyset$ and the functions f_1 and f_2 act identically on their domains (thus $V(F) = (V(F_1) \setminus \{v\}) \cup V(F_2)$). (See Figure 1 for a graph example.)

The unique substitution F (up to isomorphism) of v in F_1 by F_2 that has the smallest possible relation sets P^F ($P \in \mathcal{L}$) is called the *conservative substitution of v in F_1 by F_2* and is denoted $F_1^{v \rightarrow F_2}$. If $V(F_1) \cap V(F_2) = \emptyset$, then we can formally define $F_1^{v \rightarrow F_2}$ by

$$\begin{aligned} V(F_1^{v \rightarrow F_2}) &\stackrel{\text{def}}{=} (V(F_1) \setminus \{v\}) \cup V(F_2), \\ P^{F_1^{v \rightarrow F_2}} &\stackrel{\text{def}}{=} P^{F_2} \cup \{f_u \circ \alpha \mid \alpha \in P^{F_1} \wedge u \in V(F_2)\}, \end{aligned}$$

for every $P \in \mathcal{L}$, where $f_u: V(F_1) \rightarrow V(F_2)$ is the function that acts identically on $V(F_1) \setminus \{v\}$ and has $f_u(v) = u$.

We say that a family \mathcal{F} of \mathcal{L} -structures (up to isomorphism) is *strongly closed under substitutions* if for every $F_1, F_2 \in \mathcal{F}$, every $v \in V(F_1)$ and every substitution F of v in F_1 by F_2 , we have $F \in \mathcal{F}$. We say that \mathcal{F} is *weakly closed under substitutions* if for every $F_1, F_2 \in \mathcal{F}$, every $v \in V(F_1)$, there exists some substitution F of v in F_1 by F_2 such that $F \in \mathcal{F}$. The *strong closure under substitutions* of \mathcal{F} is the smallest family $S(\mathcal{F})$ containing \mathcal{F} that is strongly closed under substitutions.

We say that a finite \mathcal{L} -structure F is *prime*⁶ if it is not a substitution of v in F_1 by F_2 for any F_1, F_2 and $v \in V(F_1)$ with $|F_1|, |F_2| < |F|$. We say that an \mathcal{L} -structure F is *monochromatic* if for every *unary* predicate symbol $P \in \mathcal{L}$, we have $F \models \forall x \forall y, P(x) \leftrightarrow P(y)$, that is, each unary predicate is either true everywhere or true nowhere in F .

Example 3.9. To illustrate Definition 3.8, which includes subtleties that do not appear the case of graphs, suppose \mathcal{H} is a family of 3-uniform hypergraphs, with hyperedge R . Suppose our family contains F_1 , the empty hypergraph on vertices $\{v_1, v_2, v_3, v_4\}$, and F_2 is the hypergraph on vertices $\{w_1, w_2, w_3\}$ having a single hyperedge (w_1, w_2, w_3) . Suppose we substitute F_2 into F_1 by replacing the vertex v_1 by a copy of F_2 . In the hypergraph resulting from the substitution, the vertex set is $\{w_1, w_2, w_3, v_2, v_3, v_4\}$ and the requirements of the substitution are that there is a hyperedge (w_1, w_2, w_3) , and that there is no hyperedge involving exactly one of the w_i 's and exactly two of the v_j 's. However, this does not completely determine a hypergraph, since we haven't expressed an opinion about hyperedges involving two w_i 's and one v_j . Roughly speaking the *conservative substitution* is the result of saying no to all additional hyperedges, whereas saying that \mathcal{H} is *strongly closed under substitution* says any choice on these non-determined hyperedges is represented in the class.

Remark 3.10. Note that if F is a substitution of v in F_1 by F_2 , then for every unary predicate symbol $P \in \mathcal{L}$, we must have

$$\begin{aligned} (F_1 \models P(v)) &\implies (F_2 \models \forall x, P(x)), \\ (F_1 \models \neg P(v)) &\implies (F_2 \models \forall x, \neg P(x)). \end{aligned}$$

In particular, this means that any \mathcal{F} that is weakly closed under substitution can have at most one structure M_1 of size 1 (up to isomorphism), all structures F of \mathcal{F} are monochromatic and of the same “color” in the sense that for every unary predicate symbol $P \in \mathcal{L}$ and every $F \in \mathcal{F}$ with $|F| \geq 1$, we have $M_1 \models \forall x, P(x)$ if and only if $F \models \forall x, P(x)$.

Remark 3.11. Note that if F is a substitution of v in F_1 by F_2 , then $|F| \leq \max\{|F_1|, |F_2|\}$ if and only if $\min\{|F_1|, |F_2|\} \leq 1$. When F_1 has a single vertex, then $F \cong F_2$; when F_2 has a single vertex, then $F \cong F_1$; and when F_2 has no vertices (i.e., $F_2 = K_0$), then $F \cong F_1 - v$.

In particular, this means that every structure of size at most 2 is prime.

Remark 3.12. If all predicates in \mathcal{L} have arity at most 2 (which in particular covers the case of the theory of graphs), then all substitutions are conservative and the notions of weakly

⁶This should not to be confused with the notion of prime model/structure of model theory.

closed under substitutions and strongly closed under substitutions coincide. As such, in Sections 4, 5 and 6 concerning T_{Graph} , we will drop the superfluous qualifiers “weakly” and “strongly” from the terminology.

Remark 3.13. If all predicates have arity at least 3, then the notion of prime structure completely degenerates: the only prime structures are the unique structures K_0 , M_1 and M_2 of sizes 0, 1 and 2, respectively. The reason why every structure K of size at least 3 is not prime is that for any $u \in V(K)$ and $v \in V(M_2)$, K is a substitution of v in M_2 by $K - u$ since all relations involving u must involve at least two other vertices.

Remark 3.14. If T is a universal theory such that $\mathcal{M}[T]$ is weakly closed under substitution, then T is non-degenerate if and only if $\mathcal{M}_2[T] \neq \emptyset$.

Let us now prove some basic facts about substitutions and primality.

Lemma 3.15. Let \mathcal{F} be a non-empty family of \mathcal{L} -structures (up to isomorphism) that is weakly closed under substitutions. Then \mathcal{F} is closed under substructures if and only if \mathcal{F} contains the trivial structure K_0 of size 0.

Proof. Follows since a substitution of v in F by K_0 is isomorphic to $F - v$. \square

Lemma 3.16. Let F_1, F_2 be finite \mathcal{L} -structures, let $v \in V(F_1)$.

If F is a substitution of v in F_1 by F_2 and $U \subseteq V(F)$, then there exist sets $U_1 \subseteq V(F_1)$ and $U_2 \subseteq V(F_2)$ with $v \in U_1$ such that $F|_U$ is a substitution of v in $F_1|_{U_1}$ by $F_2|_{U_2}$.

Conversely, if $U_1 \subseteq V(F_1)$ and $U_2 \subseteq V(F_2)$ are such that $v \in U_1$ and F' is a substitution of v in $F_1|_{U_1}$ by $F_2|_{U_2}$, then there exist a substitution F of v in F_1 by F_2 and a set $U \subseteq V(F)$ such that $F|_U \cong F'$.

Proof. Without loss of generality, by possibly renaming vertices, we can consider only the case when F is a standard substitution. Then it is straightforward to check that for $U_1 \stackrel{\text{def}}{=} (U \cap V(F_1)) \cup \{v\}$ and $U_2 \stackrel{\text{def}}{=} U \cap V(F_2)$, we have that $F|_U$ is a substitution of v in $F_1|_{U_1}$ by $F_2|_{U_2}$.

For the second assertion, by possibly renaming vertices, we may suppose without loss of generality that F' is also a standard substitution. Then it is straightforward to see that setting $U \stackrel{\text{def}}{=} (U_1 \setminus \{v\}) \cup U_2$, there exists a standard substitution F of v in F_1 by F_2 such that $F|_U = F'$. \square

Lemma 3.17. If \mathcal{F} is a (not necessarily finite) family of finite \mathcal{L} -structures, $F \in S(\mathcal{F})$ and P is a prime substructure of F , then P is a substructure of some $F' \in \mathcal{F}$.

Proof. The proof is by induction in the minimum length ℓ of a sequence of substitutions needed to obtain F from elements of \mathcal{F} .

If $\ell = 0$, then $F \cong F'$ for some $F' \in \mathcal{F}$, so P is a substructure of F' .

If $\ell > 0$, then F is a substitution of v in M_1 by M_2 for some $M_1, M_2 \in S(\mathcal{F})$ and some $v \in V(M_1)$ such that if the minimum lengths of sequences of substitutions needed to obtain

M_1 and M_2 from elements in \mathcal{F} are ℓ_1 and ℓ_2 , respectively, then $\ell_1 + \ell_2 + 1 = \ell$, which in particular implies $\ell_1, \ell_2 < \ell$.

Without loss of generality, suppose F is a standard substitution of v in M_1 by M_2 . By Lemma 3.16, we know that there exist $U_1 \subseteq V(M_1)$ and $U_2 \subseteq V(M_2)$ with $v \in U_1$ such that P is isomorphic to some substitution of v in $M_1|_{U_1}$ by $M_2|_{U_2}$. Since P is prime, we must have $|P| \leq \max\{|U_1|, |U_2|\}$, so by Remark 3.11, either $P \cong M_1|_{U_1}$, $P \cong M_2|_{U_2}$ or $P \cong M_1|_{U_1 \setminus \{v\}}$, so P is a substructure of either M_1 or M_2 and by inductive hypothesis, it follows that P is a substructure of some $F' \in \mathcal{F}$. \square

As one might expect, prime structures play a major role in characterizing classes that are strongly closed under substitutions. This is made precise by the next two lemmas.

Lemma 3.18. *Let \mathcal{F} be a family of finite \mathcal{L} -structures (up to isomorphism) that is strongly closed under substitutions and closed under substructures and let \mathcal{P} be the set of structures in \mathcal{F} that are prime. Then $\mathcal{F} = S(\mathcal{P})$.*

Conversely, if \mathcal{P}' is a family of prime finite \mathcal{L} -structures that is closed under prime substructures and $\mathcal{F} = S(\mathcal{P}')$, then $\mathcal{P}' = \mathcal{P}$.

Proof. Let $\mathcal{F}' = S(\mathcal{P})$. It is obvious that $\mathcal{F}' \subseteq \mathcal{F}$. Suppose toward a contradiction that $\mathcal{F} \setminus \mathcal{F}' \neq \emptyset$ and let F be an \mathcal{L} -structure in $\mathcal{F} \setminus \mathcal{F}'$ of minimum size.

We claim that F is prime. Indeed, if not, then F is a substitution of some v in some F_1 by some F_2 with $|F_1|, |F_2| < |F|$. Since both F_1 and F_2 are proper substructures of F and both \mathcal{F} and \mathcal{F}' are strongly closed under substitutions and closed under substructures, this contradicts the minimality of F . Thus F is prime, so $F \in \mathcal{P}$, which contradicts $F \notin \mathcal{F}'$.

For the second assertion, if \mathcal{F} is empty, then both \mathcal{P} and \mathcal{P}' must also be empty. If \mathcal{F} is not empty, then each $P' \in \mathcal{P}'$ is in $\mathcal{F} = S(\mathcal{P})$, so Lemma 3.17 implies that P' must be a substructure of some element in \mathcal{P} , hence must also be in \mathcal{P} as it is closed under prime substructures (as \mathcal{F} is closed under substructures). Similarly, every element of \mathcal{P} must be an element of \mathcal{P}' as the latter is also closed under prime substructures. \square

Lemma 3.19. *Let T be a canonical theory in a finite relational language \mathcal{L} and let \mathcal{F} be the set of minimal \mathcal{L} -structures that are not models of T , that is, the set of all $M \in \mathcal{M}[T_{\mathcal{L}}] \setminus \mathcal{M}[T]$ such that every proper substructure of M is a model of T .*

Then $\mathcal{M}[T]$ is strongly closed under substitution if and only if \mathcal{F} contains only prime structures.

Proof. For the forward implication, note that if $M \in \mathcal{M}[T_{\mathcal{L}}] \setminus \mathcal{M}[T]$ is not prime, then it is a substitution of v in M_1 by M_2 for some $M_1, M_2 \in \mathcal{M}[T_{\mathcal{L}}]$ with $|M_1|, |M_2| < |M|$ and $v \in V(M_1)$. Since $\mathcal{M}[T]$ is strongly closed under substitutions, we must have either $M_1 \notin \mathcal{M}[T]$ or $M_2 \notin \mathcal{M}[T]$, hence $M \notin \mathcal{F}$.

For the backward implication, first note that T is a reaxiomatization of $\text{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$. Let us show that if M is a substitution of v in $M_1 \in \mathcal{M}[T]$ by $M_2 \in \mathcal{M}[T]$, then $M \in \mathcal{M}[T]$. Without loss of generality, let us assume the substitution to be standard so there is the natural identification of $V(M)$ with $(V(M_1) \setminus \{v\}) \cup V(M_2)$.

Suppose toward a contradiction that $M \notin \mathcal{M}[T]$, that is, there exists $F \in \mathcal{F}$ and $U \subseteq V(M)$ such that $M|_U \cong F$. Since $M_2 \in \mathcal{M}[T]$, we must have $U \cap V(M_1) \neq \emptyset$ (otherwise, F would be a substructure of M_2) and since $M_1 \in \mathcal{M}[T]$, we must have $|U \cap V(M_2)| \geq 2$ (otherwise, F would be a substructure of M_1). But then F is a substitution of v in $M_1|_{(U \cap V(M_1)) \cup \{v\}}$ by $M_2|_{U \cap V(M_2)}$ and since

$$\begin{aligned}|(U \cap V(M_1)) \cup \{v\}| &\leq |U| - |U \cap V(M_2)| + 1 < |U| = |F|, \\ |U \cap V(M_2)| &\leq |U| - |U \cap V(M_1)| < |U| = |F|,\end{aligned}$$

this contradicts the fact that F is prime (as it is an element of \mathcal{F}). Thus $\mathcal{M}[T]$ is strongly closed under substitutions. \square

Next we cover the notion of an almost finite family of structures and some associated notions.

Definition 3.20. We say that a family \mathcal{F} of \mathcal{L} -structures (up to isomorphism) is *almost finite* if \mathcal{F} does not contain any infinite antichain in the substructure partial order. Equivalently, \mathcal{F} is almost finite if for every infinite $\mathcal{F}' \subseteq \mathcal{F}$, there exist $F_1, F_2 \in \mathcal{F}'$ such that F_1 is a proper substructure of F_2 .

Given a family \mathcal{F} of \mathcal{L} -structures (up to isomorphism), let $\mathcal{P} \subseteq \mathcal{F}$ be the set of prime structures in \mathcal{F} and $\mathcal{P}' \subseteq \mathcal{P}$ be the set of monochromatic prime structures in \mathcal{F} . We say that \mathcal{F} is

- i. *primally finite*, if \mathcal{P} is finite,
- ii. *primally almost finite*, if \mathcal{P} is almost finite,
- iii. *monochromatically primally finite*, if \mathcal{P}' is finite,
- iv. *monochromatically primally almost finite*, if \mathcal{P}' is almost finite.

Example 3.21. An example of a family of graphs that is primally almost finite, strongly (equivalently, weakly) closed under substitutions and closed under substructures but is not primally finite is $S(\{P_n \mid n \in \mathbb{N}\})$, where P_n is the path on n vertices.

An example of a proper family of graphs that is not primally almost finite, but is strongly closed under substitutions and closed under substructures is $S(\{K_0\} \cup \{C_n \mid n \geq 5\})$, where C_n is the cycle on n vertices (it is straightforward to check that when $n \geq 5$, C_n is prime).

Another very important such example is the family $S(\{G_n \mid n \geq 6\} \cup \{K_0\})$, where G_n is the graph obtained from P_n by adding four vertices a, b, c, d and adding the edges $\{a, b\}, \{c, d\}, \{a, 2\}, \{b, 3\}, \{c, n-2\}, \{d, n-1\}$, assuming that $V(P_n) = [n]$ in the natural order of the path (see Figure 2). It is straightforward to check that each G_n is prime and that they are pairwise incomparable in the induced subgraph partial order. Note also that the paths P_n are elements of $S(\{G_n \mid n \geq 6\} \cup \{K_0\})$ as P_n is a substructure of G_n .

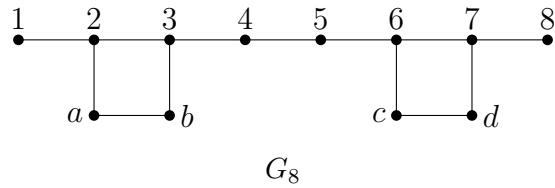
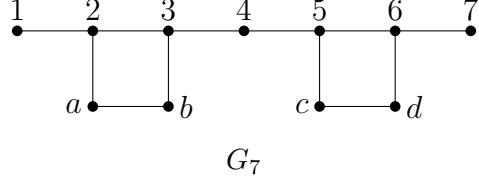
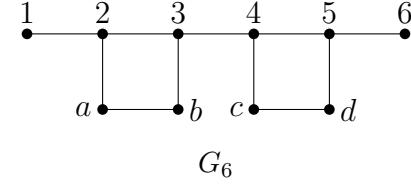


Figure 2: Prime graphs G_n of Example 3.21 that form a family that is not almost finite.

Remark 3.22. A family \mathcal{F} of the form $\mathcal{F} = S(\mathcal{P}')$ for some almost finite set of prime structures \mathcal{P}' is not necessarily primally almost finite; this is because the set \mathcal{P} of structures in \mathcal{F} that are prime is equal to the set of prime substructures of elements of \mathcal{P}' , which may be a proper superset of \mathcal{P}' (see Lemma 3.17).

As an example, consider the graphs G_n of Example 3.21 and for each $n \geq 6$, define the (prime) graph H_n as the graph obtained from the disjoint union of G_6, \dots, G_n by connecting all first vertices of the P_k inside of G_k in a clique (see Figure 3). Obviously, each H_n is an induced subgraph of H_{n+1} , so $\mathcal{P}' \stackrel{\text{def}}{=} \{K_0\} \cup \{H_n \mid n \geq 6\}$ is almost finite, but since $\{G_n \mid n \geq 6\} \subseteq S(\mathcal{P}')$, it follows that $S(\mathcal{P}')$ is not primally almost finite.

The next lemma uses the fact that the substructure partial order on finite structures is well-founded to provide a useful equivalent formulation of almost finiteness.

Lemma 3.23. *The following are equivalent for a family \mathcal{F} of finite \mathcal{L} -structures (up to isomorphism).*

- i. *The family \mathcal{F} is almost finite.*
- ii. *For every sequence $(F_n)_{n \in \mathbb{N}}$ in \mathcal{F} , there exist $n, m \in \mathbb{N}$ such that $n < m$ and F_n is a substructure of F_m .*

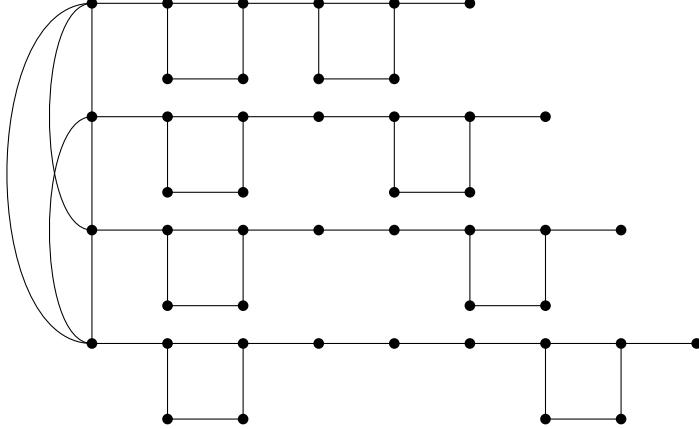


Figure 3: Prime graph H_9 of Remark 3.22. Even though the family $\{K_0\} \cup \{H_n \mid n \geq 6\}$ is almost finite, the family $S(\{K_0\} \cup \{H_n \mid n \geq 6\})$ is not primally almost finite.

Proof. We start with the implication (i) \implies (ii).

If there exist $n, m \in \mathbb{N}$ with $n < m$ and $F_n \cong F_m$, then F_n is a substructure of F_m . Suppose then that the F_n are pairwise non-isomorphic. Let I be the set of all $n \in \mathbb{N}$ such that for every $m \in \mathbb{N}$, F_m is not a proper substructure of F_n .

We claim that I is finite. Indeed, otherwise, $\mathcal{F}' \stackrel{\text{def}}{=} \{F_n \mid n \in I\}$ would be an infinite subfamily of \mathcal{F} such that for all $F, F' \in \mathcal{F}'$, F is not a proper substructure of F' .

We now construct inductively a sequence m_0, \dots, m_t as follows. Let $m_0 > \max(I)$. Given m_i , if $m_i \notin I$, then there exists m_{i+1} such that $F_{m_{i+1}}$ is a proper substructure of F_{m_i} ; otherwise, we set $t \stackrel{\text{def}}{=} i$ and stop the construction.

Since $|F_{i+1}| < |F_i|$ and all structures are finite, the construction above must stop and by a simple induction, we have that F_{m_t} is a proper substructure of F_{m_0} . Finally, since $m_t \in I$ and $m_0 > \max(I)$, we get $m_t < m_0$.

Let us now show the implication (ii) \implies (i). Let \mathcal{F}' be an infinite subfamily of \mathcal{F} and enumerate it as $(F_n)_{n \in \mathbb{N}}$ without repetitions. Then there exists $n, m \in \mathbb{N}$ such that $n < m$ and F_n is a substructure of F_m . Since $F_n \neq F_m$, it follows that F_n is a proper substructure of F_m . \square

We end this section with the following proposition that relates the notions of primally finite and primally almost finite for a family \mathcal{F} strongly closed under substitutions and closed under substructures with the number of subclasses of \mathcal{F} that are strongly closed under substitutions and closed under substructures.

Proposition 3.24. *Let \mathcal{F} be a family of finite \mathcal{L} -structures that is strongly closed under substitutions and closed under substructures, let \mathcal{S} be the set of subfamilies \mathcal{F}' of \mathcal{F} that are strongly closed under substitutions and closed under substructures. Then the following hold.*

- i. \mathcal{F} is primally finite if and only if \mathcal{S} is finite.

ii. \mathcal{F} is primally almost finite if and only if \mathcal{S} is countable.

Proof. Let \mathcal{P} be the set of prime structures of \mathcal{F} and let \mathcal{S}' be the set of subfamilies $\mathcal{P}' \subseteq \mathcal{P}$ that are closed under prime substructures. Then Lemma 3.18 gives a natural bijection between \mathcal{S} and \mathcal{S}' .

Since \mathcal{P} is finite if and only if \mathcal{S}' is finite, item (i) follows.

The same bijection between \mathcal{S} and \mathcal{S}' implies that for item (ii), it is sufficient to show that \mathcal{P} almost finite is equivalent to \mathcal{S}' countable.

Suppose first that \mathcal{P} is not almost finite, that is, \mathcal{P} contains an infinite (countable) antichain $\mathcal{A} \subseteq \mathcal{P}$. Then for each $\mathcal{A}' \subseteq \mathcal{A}$, let $\mathcal{C}_{\mathcal{A}'} \subseteq \mathcal{P}$ be the set of elements of \mathcal{P} that are substructures of some element of \mathcal{A}' . Since \mathcal{A} is an antichain, it follows that for $\mathcal{A}', \mathcal{A}'' \subseteq \mathcal{A}$ distinct, we have $\mathcal{C}_{\mathcal{A}'} \neq \mathcal{C}_{\mathcal{A}''}$, hence \mathcal{S}' has cardinality of the continuum.

Suppose now that \mathcal{S}' is uncountable. Let us define by induction in $n \in \mathbb{N}$ two sequences $(P_n)_{n \in \mathbb{N}}$ and $(\mathcal{S}'_n)_{n \in \mathbb{N}}$ with the following properties.

1. $\mathcal{S}'_n \subseteq \mathcal{S}'$ is uncountable.
2. There exists $\mathcal{C} \in \mathcal{S}'_n$ such that $P_n \in \mathcal{C}$.
3. For every $\mathcal{C}' \in \mathcal{S}'_{n+1}$, we have $P_n \notin \mathcal{C}'$.

We start by setting $\mathcal{S}'_0 \stackrel{\text{def}}{=} \mathcal{S}'$. Given \mathcal{S}'_n , let $\mathcal{P}_n \stackrel{\text{def}}{=} \bigcup_{\mathcal{C} \in \mathcal{S}'_n} \mathcal{C}$ and for each $P \in \mathcal{P}_n$, we let $\mathcal{S}'_n(P) \stackrel{\text{def}}{=} \{\mathcal{C} \in \mathcal{S}'_n \mid P \notin \mathcal{C}\}$. Since \mathcal{S}'_n is uncountable, \mathcal{P}_n is countable and $\mathcal{S}'_n \subseteq \{\mathcal{P}_n\} \cup \bigcup_{P \in \mathcal{P}_n} \mathcal{S}'_n(P)$, by the pigeonhole principle, there exists $P_n \in \mathcal{P}_n$ such that $\mathcal{S}'_n(P_n)$ is uncountable. Set $\mathcal{S}'_{n+1} \stackrel{\text{def}}{=} \mathcal{S}'_n(P_n)$ so that by definition, we have $P_n \in \mathcal{C}$ for some $\mathcal{C} \in \mathcal{S}'_n$ but $P_n \notin \mathcal{C}'$ for every $\mathcal{C}' \in \mathcal{S}'_{n+1}$. This concludes the construction.

Since each $\mathcal{C} \in \mathcal{S}'$ is closed under prime substructures, items (1), (2) and (3) together imply that for every $n < m$, P_n is not a substructure of P_m , so by Lemma 3.23, we get that \mathcal{P} is not almost finite. \square

4 Persistence in graphons

In this section we study the notion of (strongly) persistent families of graphs (Definition 4.1 below). The main objective of this section is to characterize persistence for theories of graphs in terms of closure under substitution and under induced subgraphs. We also remind the reader that in this section we drop the qualifiers “weakly” and “strongly” from “closed under substitutions” as they are superfluous for graphs (see Remark 3.12).

Definition 4.1. Let W be a graphon. The set of *positive graphs in W* is the set $Q(W) \stackrel{\text{def}}{=} \mathcal{M}[\text{Th}(W)]$ of all finite graphs G (up to isomorphism) such that $\phi_W(G) > 0$. The set of *persistently positive graphs in W* is the set $P(W) \stackrel{\text{def}}{=} \bigcap_{W'} Q(W')$, where the intersection is over all subgraphons W' of W . Clearly, $P(W)$ and $Q(W)$ depend only on the limit

$\phi_W \in \text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$. Thus, for $\phi \in \text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$, we let $P(\phi) \stackrel{\text{def}}{=} P(\phi_W)$ and $Q(\phi) \stackrel{\text{def}}{=} Q(\phi_W)$ for any graphon W such that $\phi = \phi_W$.

A graphon W (or the limit ϕ_W it represents) is called *weakly random* if $P(W) = Q(W)$. Equivalently, W is weakly random if $Q(W) = Q(W')$ for every subgraphon W' of W .

A family of graphs (up to isomorphism) \mathcal{F} is called *persistent* if there exists a graphon W such that $P(W) = \mathcal{F}$. The family \mathcal{F} is called *strongly persistent* if there exists a weakly random W such that $P(W) = \mathcal{F}$ (which must also equal $Q(W)$); in this case, we also say that W (or rather ϕ_W) is a *universal weakly random limit of \mathcal{F}* . If T is a universal theory of graphs, then we say that W is a universal weakly random limit of T if it is a universal weakly random limit of $\mathcal{M}[T]$.

Obviously, both $Q(W)$ and $P(W)$ are closed under induced subgraphs, $Q(W') \subseteq Q(W)$ whenever W' is a subgraphon of W , $P(W) \subseteq Q(W)$ and every strongly persistent family is persistent.

Example 4.2. The simplest weakly random graphons are of course the two trivial graphons, that is, the clique $W \equiv 1$ and the empty graphon $W \equiv 0$.

The next in line are the non-trivial quasirandom graphons $W_p \equiv p$ for some $p \in (0, 1)$: this is because just as the trivial graphons, the quasirandom graphons W_p also have the property that the only subgraphon of W_p is W_p , up to zero-measure change. In fact, it is an immediate consequence of the classic theory of graph quasirandomness [CGW89] that this property characterizes the quasirandom graphons (see also [SS97, SS03] for related graph quasirandomness properties); for general theories in finite relational languages, this property is called `UIInduce[1]` in [CR23] and is a strengthening of the more well-known discrepancy and clique-discrepancy properties from hypergraph quasirandomness (see [CG90, Chu90, KRS02, LM15, Tow17, AHCH⁺18]).

Other examples of weakly random graphons are the recursive blow-up of C_4 (see Proposition 5.9) and the graphon of agreements of the quasirandom permutoon (see Proposition 5.11).

The following lemma is a simple but very powerful observation about persistent families.

Lemma 4.3. *Let W be a graphon over a space $\Omega = (X, \mathcal{A}, \mu)$. Then $P(W) = \bigcap_A Q(W|_A)$, where the intersection is over all positive measure sets $A \subseteq X$.*

Proof. Since each $W|_A$ is a subgraphon of W , it is sufficient to show that if $H \in \bigcap_A Q(W|_A)$, then $H \in P(W)$. We prove this by the contra-positive: if there exists a subgraphon W' of W such that $H \notin Q(W')$, then $\phi_{W'} = \phi_{W|_f}$ for some measurable function $f: X \rightarrow [0, 1]$. Let $A \stackrel{\text{def}}{=} \{x \in X \mid f(x) > 0\}$. It is easy to see that $Q(W|_f) = Q(W|_A)$, thus $H \notin \bigcap_A Q(W|_A)$. \square

The objective of this section is to prove the following theorem that characterizes (strongly) persistent families in terms of substitutions.

Theorem 4.4. *The following are equivalent for a family \mathcal{F} of finite graphs (up to isomorphism) containing at least one graph of size at least 2.*

- i. *The family \mathcal{F} is strongly persistent.*

ii. The family \mathcal{F} is persistent.

iii. The family \mathcal{F} is closed under substitutions and induced subgraphs.

Example 4.5. As we will see, both the theory T_{Perfect} of perfect graphs (Proposition 5.10) and the theory $T_{\text{PermGraph}}$ of graphs of agreements of permutations (Proposition 5.12) have their corresponding $\mathcal{M}[T]$ closed under substitutions, hence both these theories have a universal weakly random graphon (in fact, Proposition 5.11 gives an example for the latter theory that is very different from the one produced in the proof of Theorem 4.4).

On the negative side, the theory of triangle-free graphs does not have $\mathcal{M}[T]$ closed under substitutions (as $K_2 \in \mathcal{M}[T]$ but $K_3 \cong K_2^{v \rightarrow K_2} \notin \mathcal{M}[T]$), so it does not have a universal weakly random graphon.

We will prove Theorem 4.4 through a series of lemmas. As we noted before, the implication (i) \implies (ii) is trivial. The implication (ii) \implies (iii) is a corollary of the following lemma.

Lemma 4.6. *If W is a graphon, then $P(W)$ is closed under substitutions and induced subgraphs.*

Proof. It is obvious that $K_0 \in P(W)$, so by Lemma 3.2, it is sufficient to show that $P(W)$ is closed under substitutions.

Let $F_1, F_2 \in P(W)$ and $v_0 \in V(F_1)$ and let us show that if W' is a subgraphon of W , then $t_{\text{ind}}(F_1^{v_0 \rightarrow F_2}, W') > 0$. Without loss of generality, we suppose $V(F_1) \cap V(F_2) = \emptyset$. Suppose toward a contradiction that $t_{\text{ind}}(F_1^{v_0 \rightarrow F_2}, W') = 0$. By possibly applying the Graphon Removal Lemma [Pet13, Theorem 1] to W' , we may suppose that the set $T_{\text{ind}}(F_1^{v_0 \rightarrow F_2}, W')$ of copies of $F_1^{v_0 \rightarrow F_2}$ in W' is contained in the diagonal set $\mathcal{D}_{V(F_1^{v_0 \rightarrow F_2})}$ (see (3)).

Since $F_1 \in P(W)$, we must have $t_{\text{ind}}(F_1, W') > 0$, that is, the set $T_{\text{ind}}(F_1, W')$ has positive measure. For every $(x, y) \in X^{V(F_1) \setminus \{v_0\}} \times [0, 1]^{\binom{V(F_1)}{2}}$, let

$$U_{x,y} \stackrel{\text{def}}{=} \{z \in X^{\{v_0\}} \mid ((x, z), y) \in T_{\text{ind}}(F_1, W')\}.$$

By Fubini's Theorem, there exists $(x, y) \in X^{V(F_1) \setminus \{v_0\}} \times [0, 1]^{\binom{V(F_1)}{2}}$ with all x coordinates distinct such that $U_{x,y}$ has positive measure. Since $W'|_{U_{x,y}}$ is a subgraphon of W' , hence of W , we must have $t_{\text{ind}}(F_2, W'|_{U_{x,y}}) > 0$, which implies that there exists $(z, w) \in X^{V(F_2)} \times [0, 1]^{\binom{V(F_2)}{2}}$ with all z coordinates in $U_{x,y}$, distinct and distinct from those in x such that $z \in T_{\text{ind}}(F_2, W')$. Thus, the point $(\hat{x}, \hat{y}) \in X^{V(F_1^{v_0 \rightarrow F_2})} \times [0, 1]^{\binom{V(F_1^{v_0 \rightarrow F_2})}{2}}$ defined by

$$\hat{x}_v \stackrel{\text{def}}{=} \begin{cases} x_v, & \text{if } v \in V(F_1) \setminus \{v_0\}, \\ z_v, & \text{if } v \in V(F_2), \end{cases} \quad \hat{y}_A \stackrel{\text{def}}{=} \begin{cases} y_A, & \text{if } A \subseteq V(F_1) \setminus \{v_0\}, \\ w_A, & \text{if } A \subseteq V(F_2), \\ y_{(A \cap V(F_1)) \cup \{v_0\}}, & \text{otherwise.} \end{cases}$$

is a point in $T_{\text{ind}}(F_1^{v_0 \rightarrow F_2}, W') \setminus \mathcal{D}_{V(F_1^{v_0 \rightarrow F_2})}$, a contradiction. \square

To show the final implication (iii) \implies (i) of Theorem 4.4, we will use the repeating recursive blow-up relative to an infinite sequence of graphs defined below.

Definition 4.7. Given a sequence $V = (V_\ell)_{\ell \in \mathbb{N}}$ of finite sets with $|V_\ell| \geq 2$ for every $\ell \in \mathbb{N}$, the *Cantor probability space* corresponding to V is the space $\Omega^V = (\prod_{\ell \in \mathbb{N}} V_\ell, \mathcal{A}, \nu^V)$, where \mathcal{A} is the Borel σ -algebra of the product topology on $\prod_{\ell \in \mathbb{N}} V_\ell$ and ν^V is the unique Borel measure such that $\nu^V(K_{\sigma, V}) = \prod_{\ell=0}^{t-1} |V_\ell|^{-1}$ for every $t \in \mathbb{N}$ and every $\sigma \in \prod_{\ell=0}^{t-1} V_\ell$, where $K_{\sigma, V}$ is the basic clopen set given by

$$K_{\sigma, V} \stackrel{\text{def}}{=} \left\{ \tau \in \prod_{\ell \in \mathbb{N}} V_\ell \mid \forall \ell \in \{0, \dots, t-1\}, \tau_\ell = \sigma_\ell \right\}. \quad (8)$$

Let $G = (G_m)_{m \in \mathbb{N}}$ be a sequence of finite graphs with $|G_m| \geq 2$ for every $m \in \mathbb{N}$, we let the *recursive blow-up relative to G* be the limit $\phi_G \in \text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$ defined as follows. We let $V = (V_\ell)_{\ell \in \mathbb{N}}$ be defined by $V_\ell \stackrel{\text{def}}{=} V(G_\ell)$ and define the graphon W^G over the space Ω^V by

$$W^G(x, y) \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if } x \neq y \text{ and } \{x_\ell, y_\ell\} \in E(G_\ell), \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where ℓ is the first position in which x and y differ. Finally, we define $\phi_G \stackrel{\text{def}}{=} \phi_{W^G} \in \text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$ (see Example 4.10 and Figures 4 and 5 for examples).

We let the *repeating recursive blow-up relative to G* be the limit $\phi_G^* \in \text{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$ defined as follows. For each $\ell \in \mathbb{N}$, we let

$$m_\ell \stackrel{\text{def}}{=} \max\{m \in \mathbb{N} \mid 2^m \text{ divides } \ell + 1\}, \quad (10)$$

we let $G^* \stackrel{\text{def}}{=} (G_{m_\ell})_{\ell \in \mathbb{N}}$ and we define $\phi_G^* \stackrel{\text{def}}{=} \phi_{G^*}$.

Remark 4.8. Since W^G is $\{0, 1\}$ -valued, we can interpret it as a (measurable) graph H with vertex set Ω^V and the reader familiarized with lexicographic products of graphs should note that H is simply the infinite lexicographic product of $(G_m)_{m \in \mathbb{N}}$.

Remark 4.9. The definition of the numbers m_ℓ in (10) guarantees a simple but very useful property: for every $m \in \mathbb{N}$ there exist infinitely many $\ell \in \mathbb{N}$ with $m_\ell = m$. In fact, for every $m \in \mathbb{N}$ and every $\ell \in \mathbb{N}$, there exists $\ell' \in \mathbb{N}$ with $\ell < \ell' \leq \ell + 2^m$ such that $m_{\ell'} = m$, that is, for every $m \in \mathbb{N}$, we only need to wait at most 2^m steps to see m in the sequence $(m_\ell)_{\ell \in \mathbb{N}}$ regardless of where we start.

Example 4.10. If the sequence G consists of only one graph G' repeated infinitely many times, then ϕ_G is the limit of the usual notion of recursive blow-ups of a single graph G' on progressively more and more levels.

For example, the limit ϕ_{C_4} of recursive blow-ups of C_4 used in [CM22, Definition 8.5] (see Figure 4) is obtained as ϕ_G (or ϕ_G^*) when G is the sequence that is constant equal to

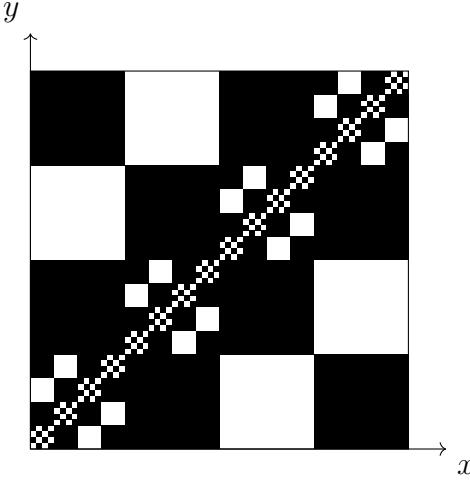


Figure 4: Approximation of a graphon W^{C_4} over $[0, 1]$ representing the limit ϕ_{C_4} of recursive blow-ups of C_4 of Example 4.10. The graphon W^{C_4} has a fractal structure, whose first 3 steps are represented in the picture.

C_4 . Alternatively, ϕ_{C_4} is also obtained as $\phi_{G'}$ for the sequence $G' \stackrel{\text{def}}{=} (K_2, \overline{K}_2, K_2, \overline{K}_2, \dots)$ that infinitely alternates between the edge graph K_2 and the non-edge graph \overline{K}_2 . Also alternatively, ϕ_{C_4} is obtained as $\phi_{G''}^*$ for the sequence $G'' = (K_2, \overline{K}_2, \overline{K}_2, \dots)$ whose first element is K_2 and all other elements are \overline{K}_2 (as $(G'')^* = G'$).

Let us show a simple structural fact about the Cantor probability space Ω^V .

Lemma 4.11. *Let $V = (V_\ell)_{\ell \in \mathbb{N}}$ be a sequence of finite sets with $|V_\ell| \geq 2$ for every $\ell \in \mathbb{N}$ and let $A \subseteq \Omega^V$ be a set with positive measure. Then for every $\epsilon > 0$, there exists $t_0 \in \mathbb{N}$ such that for every $t \geq t_0$, there exists $\sigma \in \prod_{\ell=0}^{t-1} V_\ell$ such that $\nu^V(A \cap K_{\sigma, V}) \geq (1 - \epsilon) \cdot \nu^V(K_{\sigma, V})$, where*

$$K_{\sigma, V} \stackrel{\text{def}}{=} \left\{ \tau \in \prod_{\ell \in \mathbb{N}} V_\ell \mid \forall \ell \in \{0, \dots, t-1\}, \tau_\ell = \sigma_\ell \right\}$$

is the basic clopen set defined in (8).

Proof. Let \mathcal{B} be the Boolean algebra generated by $\mathcal{C} \stackrel{\text{def}}{=} \{K_{\sigma, V} \mid t \in \mathbb{N} \wedge \sigma \in \prod_{\ell=0}^{t-1} V_\ell\}$ and note that every set in \mathcal{B} is a finite union of elements of \mathcal{C} . In fact, since for every $t \in \mathbb{N}$ and every $\sigma \in \prod_{\ell=0}^{t-1} V_\ell$, the collection $\{K_{(\sigma, v), V} \mid v \in V_t\}$ forms a partition of $K_{\sigma, V}$, it follows that for every $B \in \mathcal{B}$, there exists $t_B \in \mathbb{N}$ such that for every $t \geq t_B$, the set B can be written as the disjoint union $B = \bigcup_{\sigma \in \Sigma_{B, t}} K_{\sigma, V}$ where

$$\Sigma_{B, t} \stackrel{\text{def}}{=} \left\{ \sigma \in \prod_{\ell=0}^{t-1} V_\ell \mid K_{\sigma, V} \subseteq B \right\}.$$

Namely, we can take any representation of B as a finite union of elements of \mathcal{C} and let t_B be the maximum length of a σ used in this representation.

Since \mathcal{B} generates the σ -algebra \mathcal{A} of Ω^V , uniqueness of Carathéodory's Extension Theorem implies that for every $\delta > 0$, there exists $B \in \mathcal{B}$ such that $\nu^V(A \triangle B) \leq \delta$.

We claim that by taking $\delta \stackrel{\text{def}}{=} \nu^V(A) \cdot \epsilon / (1 + \epsilon)$ and $t_0 \stackrel{\text{def}}{=} t_B$, we get that for every $t \geq t_0$, there must exist $\sigma \in \Sigma_{B,t}$ such that $\nu^V(K_{\sigma,V} \setminus A) \leq \epsilon \cdot \nu^V(K_{\sigma,V})$. Suppose not. Then we have

$$\delta \geq \nu^V(A \triangle B) \geq \sum_{\sigma \in \Sigma_{B,t}} \nu^V(K_{\sigma,V} \setminus A) > \sum_{\sigma \in \Sigma_{B,t}} \epsilon \cdot \nu^V(K_{\sigma,V}) = \epsilon \cdot \nu^V(B) \geq \epsilon(\nu^V(A) - \delta),$$

from which we conclude that

$$\delta \cdot \frac{1 + \epsilon}{\epsilon} > \nu^V(A),$$

contradicting the definition of δ . Finally, from $\nu^V(K_{\sigma,V} \setminus A) \leq \epsilon \cdot \nu^V(K_{\sigma,V})$, we conclude that $\nu^V(A \cap K_{\sigma,V}) \geq (1 - \epsilon) \cdot \nu^V(K_{\sigma,V})$ as desired. \square

Our next objective is to show that $P(\phi_G^*)$ is precisely $S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$. We start by showing the simpler fact $\{G_m \mid m \in \mathbb{N}\} \subseteq Q(\phi_G) \subseteq S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$ in Lemma 4.12 below. Clearly this implies the same statement for ϕ_G^* .

Lemma 4.12. *Let $G = (G_m)_{m \in \mathbb{N}}$ be a sequence of finite graphs with $|G_m| \geq 2$ for every $m \in \mathbb{N}$. Then $\{G_m \mid m \in \mathbb{N}\} \subseteq Q(\phi_G) \subseteq S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$.*

Proof. To see that every G_m has positive density in ϕ_G , simply note that if we take an arbitrary $\sigma \in \prod_{\ell=0}^{m-1} V(G_\ell)$ and set $x_v \stackrel{\text{def}}{=} (\sigma, v)$ for every $v \in V(G_m)$, then x is an embedding of G_m in W^G and thus

$$t_{\text{ind}}(G_m, W^G) \geq \prod_{\ell=0}^m \frac{1}{|G_\ell|} > 0,$$

hence $\{G_m \mid m \in \mathbb{N}\} \subseteq Q(\phi_G)$.

To show that $Q(\phi_G) \subseteq S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$, we will prove a slightly stronger result: let us show that if H is a finite graph such that $T_{\text{ind}}(H, W^G) \not\subseteq \mathcal{D}_{V(H)}$, then $H \in S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$. The proof is by induction on $|H|$.

The first two base cases are when $|H| \leq 1$ (i.e., $H \in \{K_0, K_1\}$), in which case trivially $H \in S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$.

The next base cases are when $|H| \geq 2$ and H is prime. In this case, we show that H must be an induced subgraph of G_m for some $m \in \mathbb{N}$.

Recall that since W^G is $\{0, 1\}$ -valued, the set $T_{\text{ind}}(H, W^G) \setminus \mathcal{D}_{V(H)}$ is alternatively described as the set of pairs $(x, y) \in (\Omega^V)^{V(H)} \times [0, 1]^{\binom{V(H)}{2}}$ such that x is an embedding of H in W^G .

Fix then one such point (x, y) and let $t \in \mathbb{N}$ be the length of the longest string $\sigma \in \prod_{\ell=0}^{t-1} V(G_\ell)$ that is common to all coordinates of x , that is, for every $v \in V(H)$, we have $x_v|_{\{0, \dots, t-1\}} = \sigma$ and there exist $v, w \in V(H)$ such that $(x_v)_t \neq (x_w)_t$.

For each $i \in V(G_t)$, let $U_i \stackrel{\text{def}}{=} \{v \in V(H) \mid (x_v)_t = i\}$ and let $H_i \stackrel{\text{def}}{=} H|_{U_i}$. Let also $I \stackrel{\text{def}}{=} \{i \in V(G_t) \mid U_i \neq \emptyset\}$. The structure of W^G implies that H is obtained from $G_t|_I$ by substituting each $i \in I$ by H_i . Since $|H_i| < |H|$ for every $i \in I$ and H is prime, it follows that $|G_t|_I| = |H|$ and $|U_i| = 1$ for every $i \in I$, that is, the unique function $\beta: V(H) \rightarrow V(G_t)$ such that $v \in U_{\beta(v)}$ is an embedding of H in G_t . Thus H is an induced subgraph of G_t .

We now consider the inductive step when H is not prime. Then H is of the form $F_1^{v \rightarrow F_2}$ for some graphs F_1, F_2 and $v \in V(F_1)$ with $|F_1|, |F_2| < |H|$. By inductive hypothesis, we have $F_1, F_2 \in S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$ and since this set is closed under substitutions we get $H \in S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$. \square

Lemma 4.13. *Let $G = (G_m)_{m \in \mathbb{N}}$ be a sequence of finite graphs with $|G_m| \geq 2$ for every $m \in \mathbb{N}$. Then $P(\phi_G^*) = Q(\phi_G^*) = S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$.*

Proof. By Lemma 4.12, we know that $Q(\phi_G^*) \subseteq S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$ and since $P(\phi_G^*) \subseteq Q(\phi_G^*)$, it is sufficient to prove that $S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\}) \subseteq P(\phi_G^*)$.

Let $H \in S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$ and let us show that $H \in P(\phi_G^*)$ by induction on $|H|$.

The base case is when H is a prime graph. By Lemma 3.4, we know there exists $\hat{m} \in \mathbb{N}$ such that H is an induced subgraph of $G_{\hat{m}}$, so it is sufficient to show that $G_{\hat{m}} \in P(\phi_G^*)$. In turn, by Lemma 4.3, it is sufficient to show that for every positive measure set $A \subseteq \Omega^V$, the graph $G_{\hat{m}}$ has positive density in the subgraphon $W^{G^*}|_A$.

Let ϵ be any positive number with $\epsilon < 1/|G_{\hat{m}}|$. By Lemma 4.11, there exists $t_0 \in \mathbb{N}$ such that for every $t \geq t_0$, there exists $\sigma^t \in \prod_{\ell=0}^{t-1} V(G_{m_\ell})$ such that $\nu^V(A \cap K_{\sigma^t, V}) \geq (1 - \epsilon) \cdot \nu^V(K_{\sigma^t, V})$.

Let $t \in \mathbb{N}$ be such that $t_0 < t \leq t_0 + 2^{\hat{m}}$ and $m_t = \hat{m}$ as provided by Remark 4.9. Let also

$$T \stackrel{\text{def}}{=} \left\{ \tau \in \prod_{\ell=0}^t V(G_{m_\ell}) \mid \tau|_{\{0, \dots, t-1\}} = \sigma^t \right\}.$$

Since $\{K_{\tau, V} \mid \tau \in T\}$ partitions $K_{\sigma^t, V}$ into $|T| = |G_{m_t}| = |G_{\hat{m}}|$ parts of equal measure, it follows that for every $\tau \in T$, we have

$$\nu^V(A \cap K_{\tau, V}) \geq \left(1 - \epsilon - \frac{|G_{\hat{m}}| - 1}{|G_{\hat{m}}|}\right) \nu^V(K_{\sigma^t, V}) > 0.$$

However, the definition of W^{G^*} implies that if we pick $x_v \in K_{(\sigma^t, v), V}$ for each $v \in V(G_{\hat{m}})$ (and pick any $y \in [0, 1]^{\binom{|V(G_{\hat{m}})|}{2}}$), then we get a copy of $G_{\hat{m}}$ in W^{G^*} and since for every $v \in G_{\hat{m}}$, we have $\nu^V(A \cap K_{(\sigma^t, v), V}) > 0$ (as $(\sigma^t, v) \in T$), it follows that $G_{\hat{m}}$ has positive density in $W^{G^*}|_A$.

For the inductive step when H is not prime, we must have $H = F_1^{v \rightarrow F_2}$ for some graphs F_1, F_2 and some $v \in V(F_1)$ with $|F_1|, |F_2| < |H|$. Since F_1 and F_2 are in $P(\phi_G^*)$ by inductive hypothesis and $P(\phi_G^*)$ is closed under substitutions by Lemma 4.6, it follows that $H \in P(\phi_G^*)$. \square

We can finally show Theorem 4.4 that says that a family \mathcal{F} of graphs with at least one graph of size at least 2 is strongly persistent if and only if it is persistent if and only if it is closed under substitutions and under induced subgraphs.

Proof of Theorem 4.4. The implication (i) \implies (ii) is trivial: every strongly persistent family is obviously persistent.

For the implication (ii) \implies (iii), if $\mathcal{F} = P(W)$ for some graphon W , then Lemma 4.6 implies that it is closed under substitutions and induced subgraphs.

For the final implication (iii) \implies (i), by Lemma 3.5, we have $\mathcal{F} = S(\mathcal{P})$, where \mathcal{P} is the set of graphs in \mathcal{F} that are prime. Since \mathcal{F} contains at least one graph of size at least 2, \mathcal{P} must also contain one such graph (since $S(\{K_0, K_1\}) = \{K_0, K_1\}$).

Let $G = (G_m)_{m \in \mathbb{N}}$ be an enumeration of all graphs in \mathcal{P} of size at least 2 (potentially with repetitions if \mathcal{P} is finite). Note that since $\mathcal{F} = S(\mathcal{P})$ is closed under induced subgraphs, it follows that $\mathcal{F} = S(\{K_0\} \cup \{G_m \mid m \in \mathbb{N}\})$.

By Lemma 4.13, the repeating recursive blow-up ϕ_G^* relative to G satisfies $P(\phi_G^*) = Q(\phi_G^*) = \mathcal{F}$, hence \mathcal{F} is strongly persistent. \square

5 Weak randomness in graphons

Recall that Theorem 4.4 characterizes all universal theories of graphs that contain a universal weakly random graphon. In this section, we study a related natural question (see Definition 5.1 below): when does every graphon of a universal theory of graphs contain some weakly random subgraphon? As mentioned in the introduction, this property is a generalization of AEHP (see Definition 2.9). We also remind the reader that in this section we drop the qualifiers “weakly” and “strongly” from “closed under substitutions” as they are superfluous for graphs (see Remark 3.12).

Definition 5.1. We say that a universal theory T of graphs has the *weakly random Erdős–Hajnal property* (abbreviated $T \in \text{WR}$) if every limit W of T contains a weakly random subgraphon.

Remark 5.2. Since trivial graphons are weakly random, we obviously have $\text{AEHP} \subseteq \text{WR}$. Furthermore, even though it is also natural to ask what is the class of theories that have *some* weakly random limit, it is clear that this is precisely the set of non-degenerate universal theories. This is because Ramsey’s Theorem implies that any non-degenerate theory T of graphs must either contain arbitrarily large cliques, in which case $W \equiv 1$ is a limit of T , or contain arbitrarily large anti-cliques, in which case $W \equiv 0$ is a limit of T .

Similarly to Lemma 4.3, the following lemma is a simple but very powerful observation about weakly random subgraphons.

Lemma 5.3. *A graphon W over a space $\Omega = (X, \mathcal{A}, \mu)$ contains a weakly random subgraphon W' if and only if there exists a positive measure set $A \subseteq X$ such that $W|_A$ is weakly random.*

Proof. The backward implication follows because $W|_A$ is a subgraphon of W .

For the forward implication, we know that $\phi_{W'} = \phi_{W|_f}$ for some measurable function $f: X \rightarrow [0, 1]$. Let $A \stackrel{\text{def}}{=} \{x \in X \mid f(x) > 0\}$. We claim that $W|_A$ is weakly random. Indeed, this follows from Lemma 4.3 and since for every $B \subseteq X$ of positive μ_f -measure, we have $Q(W|_A|_B) = Q(W|_f|_B)$. \square

Our next main objective is to characterize the class WR under the assumption that the set of graphs of the theory is closed under substitutions.

Theorem 5.4. *Let T be a universal theory of graphs such that $\mathcal{M}[T]$ is closed under substitutions. Then $T \in \text{WR}$ if and only if $\mathcal{M}[T]$ is primally almost finite.*

Before we prove Theorem 5.4 above, let us observe a simple corollary of it.

Corollary 5.5. *There exists a universal theory of graphs T with $\mathcal{M}[T] \subsetneq \mathcal{M}[T_{\text{Graph}}]$ and $T \notin \text{WR}$.*

Proof. The family $\{C_n \mid n \geq 5\}$ of cycles of length at least 5 is a family of prime graphs that is not almost finite (see Example 3.21). Let $\mathcal{F} = S(\{K_0\} \cup \{C_n \mid n \geq 5\})$ and since \mathcal{F} is closed under substitutions and induced subgraphs but is not primally almost finite, the universal theory $T \stackrel{\text{def}}{=} \text{Th}(\mathcal{F})$ with $\mathcal{M}[\text{Th}(\mathcal{F})] = \mathcal{F}$ satisfies $T \notin \text{WR}$ by Theorem 5.4. It is also easy to see that $\mathcal{M}[T] \subsetneq \mathcal{M}[T_{\text{Graph}}]$, as for example the prime graph G_6 of Example 3.21 is not in $\mathcal{M}[T]$ by Lemma 3.4. \square

We start by proving the easier direction of Theorem 5.4 in the lemma below. In fact, for this direction, we do not even need $\mathcal{M}[T]$ to be closed under substitutions.

Lemma 5.6. *If T is a universal theory of graphs such that $\mathcal{M}[T]$ is primally almost finite, then $T \in \text{WR}$.*

Proof. We prove the lemma by its contra-positive. Suppose $T \notin \text{WR}$ and let us show that the set \mathcal{P} of graphs of T that are prime is not almost finite. By Lemma 3.7, it is sufficient to construct a sequence $(R_n)_{n \in \mathbb{N}}$ in \mathcal{P} such that for every $n, m \in \mathbb{N}$, if $n < m$, then R_n is not an induced subgraph of R_m .

Since $T \notin \text{WR}$, there must exist a limit $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ of T that does not contain any weakly random sub-object.

We now construct sequences $(\phi_n)_{n \in \mathbb{N}}$ of sub-objects of ϕ and $(R_n)_{n \in \mathbb{N}}$ of prime graphs in $\mathcal{M}[T]$ satisfying:

- i. For every $n \in \mathbb{N}$, ϕ_{n+1} is a sub-object of ϕ_n .

- ii. For every $n \in \mathbb{N}$, $R_n \in Q(\phi_n) \setminus Q(\phi_{n+1})$.

We construct these sequences inductively as follows.

1. Set $\phi_0 \stackrel{\text{def}}{=} \phi$.
2. Given $\phi_n \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$, since ϕ_n is a sub-object of ϕ , we know that ϕ_n is not weakly random, so there exists $G_n \in Q(\phi_n) \setminus P(\phi_n)$. Let \mathcal{P}_n be the set of induced subgraphs of G_n that are prime. Since by Lemma 4.6, $P(\phi_n)$ is closed under substitutions and $G_n \in S(\mathcal{P}_n)$, there exists $R_n \in \mathcal{P}_n \setminus P(\phi_n)$ and since $Q(\phi_n)$ is closed under induced subgraphs, we get $R_n \in Q(\phi_n) \setminus P(\phi_n)$. From the definition of $P(\phi_n)$, it follows that there exists a sub-object ϕ_{n+1} of ϕ_n (hence ϕ_{n+1} is also a sub-object of ϕ) such that $R_n \in Q(\phi_n) \setminus Q(\phi_{n+1})$.

Let now $n, m \in \mathbb{N}$ be such that $n < m$. By induction, we know that ϕ_m is a sub-object of ϕ_{n+1} , so $Q(\phi_m) \subseteq Q(\phi_{n+1})$, which in turn implies that $R_n \in Q(\phi_n) \setminus Q(\phi_m)$. Since $R_m \in Q(\phi_m)$ and $Q(\phi_m)$ is closed under induced subgraphs, it follows that R_n is not an induced subgraph of R_m , concluding the proof. \square

For the other side of the characterization of WR, the proposition below shows that under appropriate hypotheses, the recursive blow-up ϕ_R of Definition 4.7 is a graphon without any weakly random subgraphon.

Proposition 5.7. *Let $R = (R_n)_{n \in \mathbb{N}}$ be a sequence of prime graphs of size at least 2 such that for each $n \in \mathbb{N}$, there exist at most finitely many $m \in \mathbb{N}$ such that R_n is an induced subgraph of R_m . Suppose also that $\prod_{n \in \mathbb{N}} (1 - 1/|R_n|) = 0$. Then ϕ_R does not contain any weakly random sub-object.*

Let us first give some intuition on the proof of Proposition 5.7. First, note that since all R_n are prime graphs and ϕ_R is obtained via a limit of recursive blow-ups, which themselves are obtained from the R_n via substitutions, it follows that copies of R_n in ϕ_R need to correspond to copies of R_n inside some R_m . The condition that each R_n is contained in at most finitely many R_m then ensures that the restriction of W^R to basic clopen sets $K_{\sigma, V}$ (see (8)) with $|\sigma|$ large enough do not have any copies of R_n . Thus, for every positive measure set A , there is some $K_{\sigma, V}$ such that $R_n \notin Q(W^R|_{A \cap K_{\sigma, V}})$ and $A \cap K_{\sigma, V}$ has positive measure. However, to use this fact show that ϕ_R does not contain any weakly random sub-object, we need to also ensure that every positive measure set A contains at least one R_n (with n depending on A), so that we conclude that $Q(W^R|_A) \neq P(W^R|_A)$ since the above argument gives $R_n \in Q(W^R|_A) \setminus Q(W^R|_{A \cap K_{\sigma, V}})$. This is where the condition $\prod_{n \in \mathbb{N}} (1 - 1/|R_n|) = 0$ comes in: we will show that any set A avoiding all R_n has measure at most $\prod_{n \in \mathbb{N}} (1 - 1/|R_n|)$.

Proof. For every $t \in \mathbb{N}$, let R^t be the shifted sequence $(R_{n+t})_{n \in \mathbb{N}}$.

Also, for each $t \in \mathbb{N}$, let m_t be the maximum $m \in \mathbb{N}$ such that R_t is an induced subgraph of R_m . Note that for every $t \in \mathbb{N}$, by Lemmas 3.3 and 4.12, we have $R_t \in Q(\phi_{R^t}) \setminus Q(\phi_{R^{m_t+1}})$ since R_t is prime and is not an induced subgraph of any $R_{t'}$ with $t' > m_t$ ⁷.

⁷In fact, since $\phi_{R^{m_t+1}}$ is a sub-object of ϕ_{R^t} , we have $Q(\phi_{R^{m_t+1}}) \subsetneq Q(\phi_{R^t})$.

To show that ϕ_R does not contain a weakly random sub-object, by Lemma 5.3, it is sufficient to show that for every positive measure set $A \subseteq \Omega^V$, the subgraphon $W^R|_A$ is not weakly random.

We claim that there exists $t \in \mathbb{N}$ such that $R_t \in Q(W^R|_A)$. Suppose not, let $n \in \mathbb{N}$ be large enough so that $\prod_{\ell=0}^{n-1} (1 - 1/|R_\ell|) < \nu^V(A)$ (recall from Definition 4.7 that ν^V is the measure in the underlying space of W^R) and consider the set

$$\Sigma \stackrel{\text{def}}{=} \left\{ \sigma \in \prod_{\ell=0}^{n-1} V(R_\ell) \mid \nu^V(A \cap K_{\sigma,V}) > 0 \right\}.$$

We claim that for every $m \in \{0, \dots, n-2\}$ and every $\tau \in \prod_{\ell=0}^{m-1} V(R_\ell)$, there exists $v_\tau \in V(R_m)$ such that (τ, v_τ) is not a prefix of any element of Σ . Indeed, otherwise, since mapping each $v \in V(R_m)$ to an element of $K_{(\tau, v), V}$ gives an embedding of R_m in W^R , we would get

$$t_{\text{ind}}(R_m, W^R|_A) \geq \prod_{v \in V(R_m)} \frac{\nu^V(A \cap K_{(\tau, v), V})}{\nu^V(A)} > 0.$$

Thus, the existence of v_τ is proved.

Let then Σ^* be the set of $\sigma \in \prod_{\ell=0}^{n-1} V(R_\ell)$ such that for every $m \in \{0, \dots, n-2\}$, we have $v_{\sigma|_{\{0, \dots, m-1\}}} \neq \sigma_m$. Our last claim says that $\Sigma \subseteq \Sigma^*$. Now it is easy to see that

$$\nu^V(A) = \sum_{\sigma \in \Sigma} \nu^V(A \cap K_{\sigma,V}) \leq \sum_{\sigma \in \Sigma^*} \nu^V(K_{\sigma,V}) = \prod_{\ell=0}^{n-1} \left(1 - \frac{1}{|R_\ell|}\right) < \nu^V(A),$$

a contradiction. This concludes the proof that there exists $t \in \mathbb{N}$ such that $R_t \in Q(W^R|_A)$.

We will now show that $W^R|_A$ is not weakly random by showing that there exists a sub-object of $W^R|_A$ in which R_t has density zero (so that we conclude $P(W^R|_A) \subsetneq Q(W^R|_A)$ as R_t is in the latter set but not in the former set).

Since $\{K_{\sigma,V} \mid \sigma \in \prod_{\ell=0}^{m_t} V(R_\ell)\}$ partitions the space Ω^V , there must exist $\sigma \in \prod_{\ell=0}^{m_t} V(R_\ell)$ such that $\nu^V(A \cap K_{\sigma,V}) > 0$ but note that $\phi_{W^R|_{K_{\sigma,V}}} = \phi_{R^{m_t+1}}$ and since $Q(W|_{A \cap K_{\sigma,V}}) \subseteq Q(W|_{K_{\sigma,V}}) = Q(\phi_{R^{m_t+1}})$ it follows that $R_t \notin Q(W|_{A \cap K_{\sigma,V}})$ as desired.

Therefore ϕ_R does not contain any weakly random sub-object. \square

Remark 5.8. The product condition $\prod_{n \in \mathbb{N}} (1 - 1/|R_n|) = 0$ in Proposition 5.7 may seem very unnatural at first. However, it is easy to see that it is necessary for ϕ_R to not contain any weakly random sub-object: for example, consider the limit ϕ_R for the sequence $R \stackrel{\text{def}}{=} (C_{n^2+5})_{n \in \mathbb{N}}$ (see Figure 5) and fixing $v_n \in V(C_{n^2+5})$ for each $n \in \mathbb{N}$, let

$$A \stackrel{\text{def}}{=} \prod_{n \in \mathbb{N}} (V(C_{n^2+5}) \setminus \{v_n\}).$$

Note that $\nu^V(A) = \prod_{n \in \mathbb{N}} (1 - 1/|C_{n^2+5}|) > 0$. On the other hand, since $C_{n^2+5} - v_n \cong P_{n^2+4}$ is the path with $n^2 + 4$ vertices, it is obvious that $W^R|_A$ represents the same limit as $W^{R'}$

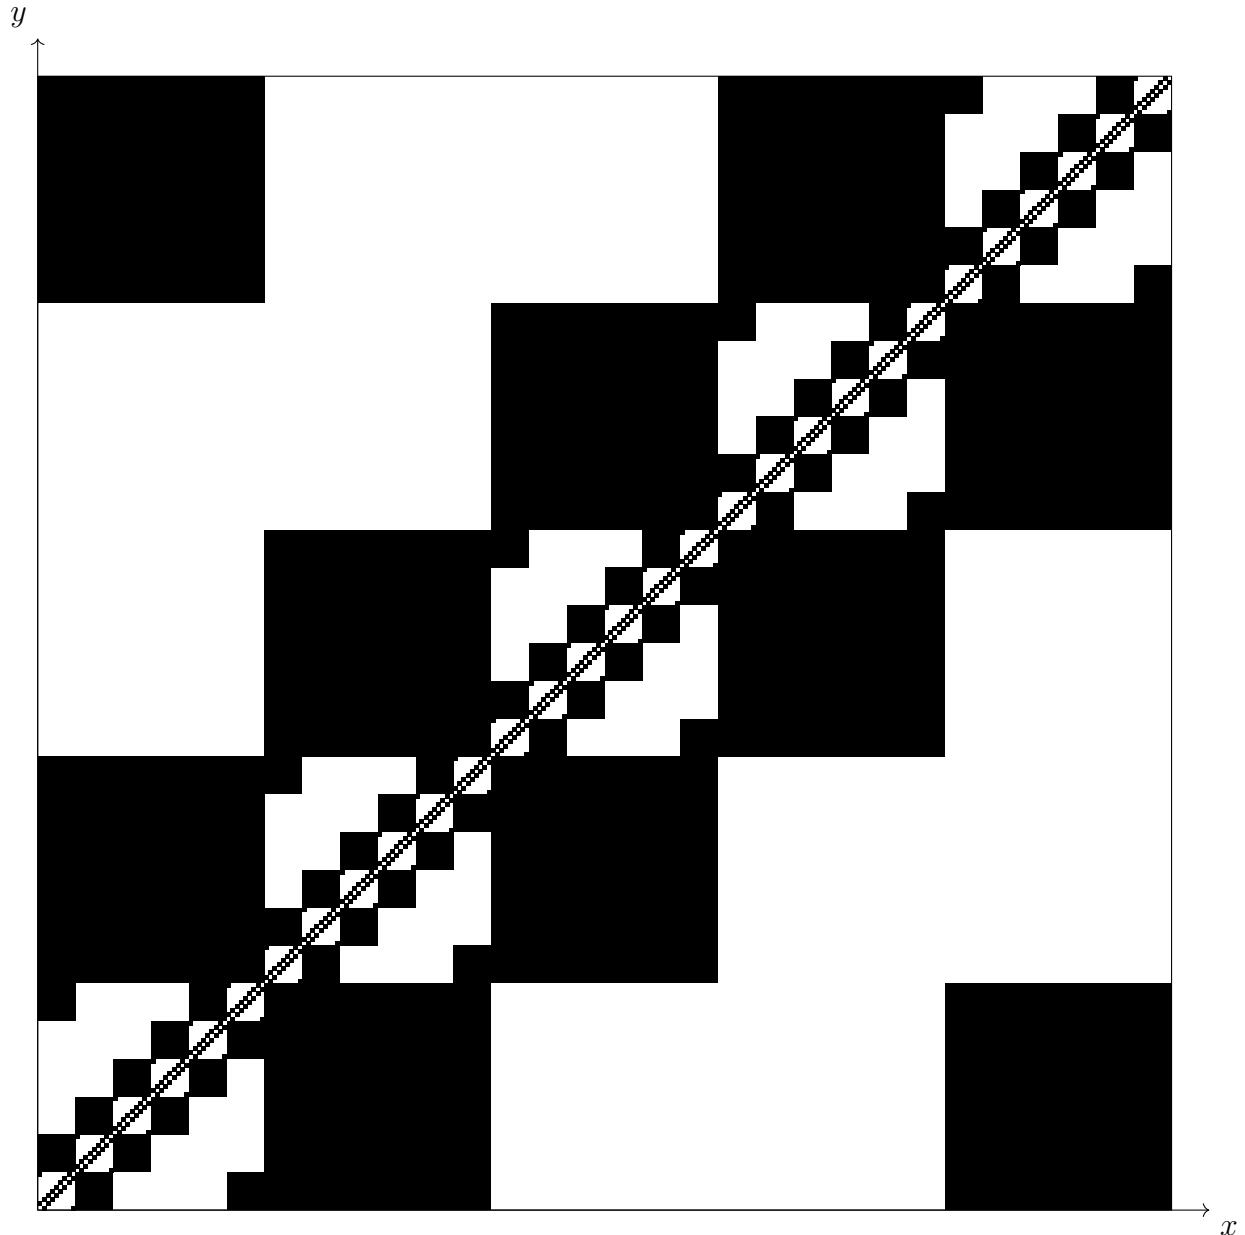


Figure 5: Approximation of a graphon W over $[0, 1]$ representing the limit ϕ_R of recursive blow-ups corresponding to the sequence $R \stackrel{\text{def}}{=} (C_{n^2+5})_{n \in \mathbb{N}}$ of Remark 5.8. The graphon W has a fractal structure, whose first 3 steps are represented in the picture.

for the sequence $R' = (P_{n^2+4})_{n \in \mathbb{N}}$. In turn, by Lemma 4.12, we have

$$Q(W^{R'}) \subseteq S(\{K_0\} \cup \{P_{n^2+4} \mid n \in \mathbb{N}\}) = S(\{P_n \mid n \in \mathbb{N}\})$$

and since the family above is primally almost finite, Lemma 5.6 implies that $W^{R'}$ contains a weakly random subgraphon, hence so does $W^{R'}$. In fact, with a bit more effort, one can also show that $W^{R'}$ itself is already weakly random, but we omit this proof. We will also see in Proposition 5.13 that not only does ϕ_R contain a weakly random sub-object, but we also have $\text{Th}(\phi_R) \in \text{WR}$.

We can now prove Theorem 5.4 that says that a universal theory of graphs T with $\mathcal{M}[T]$ closed under substitutions and with $\mathcal{M}_2[T]$ non-empty is in WR if and only if $\mathcal{M}[T]$ is primally almost finite.

Proof of Theorem 5.4. For the backward direction, if $\mathcal{M}[T]$ is primally almost finite, then by Lemma 5.6, we have $T \in \text{WR}$.

We prove the forward direction by the contra-positive: suppose $\mathcal{M}[T]$ is not primally almost finite, so there exists an infinite antichain $\{G_n \mid n \in \mathbb{N}\}$ of prime graphs of T , and without loss of generality, suppose every G_n has size at least 2.

For each $n \in \mathbb{N}$, let $r_n \in \mathbb{N}_+$ be large enough so that $(1 - 1/|G_n|)^{r_n} \leq 1/2$ and for each $\ell \in \mathbb{N}$, let $R_\ell \stackrel{\text{def}}{=} G_n$ for the unique $n \in \mathbb{N}$ such that $\sum_{m=0}^{n-1} r_m \leq \ell < \sum_{m=0}^n r_m$. Clearly, for each $\ell \in \mathbb{N}$, there exist exactly r_ℓ values of $t \in \mathbb{N}$ such that R_ℓ is an induced subgraph of R_t . On the other hand, we have

$$\prod_{\ell \in \mathbb{N}} \left(1 - \frac{1}{|R_\ell|}\right) = \prod_{n \in \mathbb{N}} \left(1 - \frac{1}{|G_n|}\right)^{r_n} \leq \prod_{n \in \mathbb{N}} \frac{1}{2} = 0.$$

By Proposition 5.7, we know that ϕ_R does not contain any weakly random sub-object and by Lemma 4.12, we know that $Q(\phi_R) \subseteq S(\{R_\ell \mid \ell \in \mathbb{N}\}) \subseteq \mathcal{M}[T]$, so ϕ_R is a limit of T without any weakly random sub-object. \square

We conclude this section with some natural examples of universal theories in WR and not in WR . We start by showing that the universal theory of induced subgraphs of recursive blow-ups of C_4 studied in [CM22, §8] (see Example 4.10 and Figure 4) is the simplest example in $\text{WR} \setminus \text{AEHP}$.

Proposition 5.9. *The limit recursive blow-up ϕ_{C_4} of C_4 is weakly random. In particular, the theory T of induced subgraphs of the recursive blow-ups of C_4 satisfies $T \in \text{WR} \setminus \text{AEHP}$.*

Proof. Recall from Example 4.10 that the limit ϕ_{C_4} recursive blow-up of C_4 can be viewed as the repeating recursive blow-up $\phi_{G''}^*$ for the sequence $G'' = (K_2, \overline{K}_2, \overline{K}_2, \dots)$ whose first element is K_2 and all other elements are \overline{K}_2 .

There are two ways of seeing that ϕ_{C_4} is weakly random. The first is using Lemma 4.13 to conclude that $P(\phi_{C_4}) = Q(\phi_{C_4}) = S(\{K_0, K_2, \overline{K}_2\})$. Alternatively, the result follows directly from the results of [CM22] and Lemma 4.6: by [CM22, Lemma 8.7], we know

that $Q(\phi_{C_4}) \subseteq S(\{K_0, K_2, \bar{K}_2\})$, so Lemma 4.6 implies that $P(\phi_{C_4})$ can only be one of $S(\{K_0, K_2\})$, $S(\{K_0, \bar{K}_2\})$ or $S(\{K_0, K_2, \bar{K}_2\})$ and since by [CM22, Lemma 8.8] does not contain trivial subgraphons, the first two cases are ruled out, so $P(\phi_{C_4}) = Q(\phi_{C_4}) = S(\{K_0, K_2, \bar{K}_2\})$.

Since the family of induced subgraphs of recursive blow-ups of C_4 is precisely the family $S(\{K_0, K_2, \bar{K}_2\})$, which is primally finite, the fact that $T \in \text{WR}$ follows from Theorem 5.4. On the other hand, ϕ_{C_4} does not contain trivial subgraphons (this follows directly from [CM22, Lemma 8.8] or alternatively from the fact that a trivial subgraphon W must have $Q(W)$ either equal to $S(\{K_0, K_2\})$ or $S(\{K_0, \bar{K}_2\})$), hence $T \notin \text{AEHP}$. \square

Proposition 5.10. *The theory T_{Perfect} of perfect graphs is not in WR . Furthermore, the set $\mathcal{M}[T_{\text{Perfect}}]$ is closed under substitutions.*

Proof. We first show that $\mathcal{M}[T_{\text{Perfect}}]$ is closed under substitutions. By the Strong Perfect Graph Theorem [CRST06], we know that a graph G is perfect if and only if both G and its complement \bar{G} do not contain any induced odd-cycle of length at least 5.

Let us show that if F_1, F_2 are perfect graphs and $v \in V(F_1)$, then $F_1^{v \rightarrow F_2}$ is also a perfect graph. Since $\overline{F_1^{v \rightarrow F_2}} \cong (\overline{F_1})^{v \rightarrow \overline{F_2}}$, it is sufficient to show that $F_1^{v \rightarrow F_2}$ does not contain any induced odd-cycles of length at least 5.

Without loss of generality, let us suppose $V(F_1) \cap V(F_2) = \emptyset$. Suppose toward a contradiction that $v_1, \dots, v_{2\ell+1}$ forms an induced odd-cycle of $F_1^{v \rightarrow F_2}$ with $\ell \geq 2$. Since both F_1 and F_2 are perfect, this odd-cycle must contain both vertices of F_1 (that are not v) and F_2 . Without loss of generality, suppose $v_i \in V(F_2)$ for every $i \in [k]$ for some $k \in [2\ell+1]$ and $v_{k+1} \in V(F_1)$. Since $v_{k+1} \in V(F_1)$ is adjacent to $v_k \in V(F_2)$, it follows from the structure of $F_1^{v \rightarrow F_2}$ that v_{k+1} is adjacent to all of v_1, \dots, v_k , but since the cycle is induced, this can only happen if $k = 2$ and $2\ell+1 = 3$, a contradiction. Therefore, $\mathcal{M}[T_{\text{Perfect}}]$ is closed under substitutions.

By Theorem 5.4, to show that $T_{\text{Perfect}} \notin \text{WR}$, it is sufficient to show that $\mathcal{M}[T_{\text{Perfect}}]$ is not primally almost finite. But recall that the family of graphs $\{G_n \mid n \geq 6\}$ of Example 3.21 is a family of prime graphs that is not almost finite and since these graphs are bipartite, they are also perfect. \square

Finally, we consider the theory $T_{\text{PermGraph}} \stackrel{\text{def}}{=} I(T_{\text{Perm}})$ of graphs of agreements of permutations, where $I: T_{\text{Graph}} \rightsquigarrow T_{\text{Perm}}$ is given by

$$I(E)(x, y) \stackrel{\text{def}}{=} (x \neq y \wedge (x \prec_1 y \leftrightarrow x \prec_2 y)),$$

The next proposition provides a natural universal weakly random limit of $T_{\text{PermGraph}}$ as the graphon of agreements of the quasirandom permutoon (see Figure 6). However, we defer its proof to Section 7 as it will follow as an easy consequence of naturality of weak randomness (Proposition 7.10(iii)) and the fact that the quasirandom permutoon is a universal weakly random limit of T_{Perm} (Proposition 7.15).

Proposition 5.11. *The graphon W over $[0, 1]^2$ of agreements of the quasirandom permuton given by*

$$W(x, y) \stackrel{\text{def}}{=} \mathbb{1}[\pi_1(x) < \pi_1(y) \leftrightarrow \pi_2(x) < \pi_2(y)],$$

where $\pi_i: [0, 1]^2 \rightarrow [0, 1]$ is the projection onto the i th coordinate, is a universal weakly random limit of $T_{\text{PermGraph}}$.

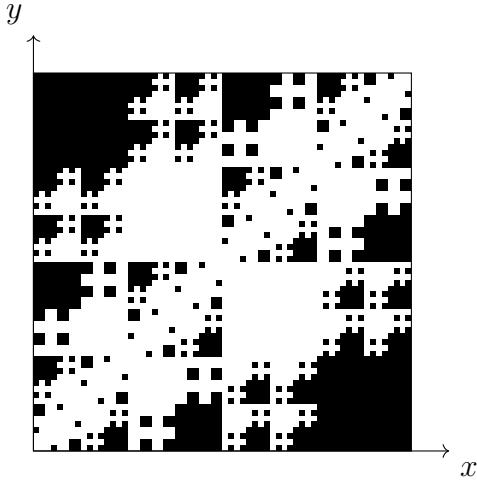


Figure 6: Approximation of a graphon W' over $[0, 1]$ representing the same limit as the graphon W of agreements of the quasirandom permuton of Proposition 5.11. The graphon W' is indirectly defined by $W'(F(x), F(y)) = W(x, y)$ for the standard measure-isomorphism F from $[0, 1]$ to $[0, 1]^2$ that maps the point $(x, y) \in [0, 1]^2$ to $\sum_{i \in \mathbb{N}_+} 4^{-i} \cdot (2 \cdot x_i + y_i)$, where the binary expansions of x and y are $0.x_1x_2\cdots$ and $0.y_1y_2\cdots$, respectively. The graphon W' has a fractal structure, whose first 3 steps are represented in the picture.

Proposition 5.12. *The theory $T_{\text{PermGraph}}$ of graphs of agreements of permutations is not in WR. Furthermore, $\mathcal{M}[T_{\text{PermGraph}}]$ is closed under substitutions.*

Proof. First let us prove that $\mathcal{M}[T_{\text{PermGraph}}]$ is closed under substitutions. Let $F, G \in \mathcal{M}[T_{\text{PermGraph}}]$, let $v \in V(F)$ and without loss of generality, suppose $V(F) = [n]$ and $V(G) = [m]$ for some $n, m \in \mathbb{N}$ and $\sigma \in S_n$ and $\tau \in S_m$ are permutations representing F and G with $\{i, j\} \in E(F)$ if and only if $i < j \leftrightarrow \sigma(i) \leq \sigma(j)$ and analogously for G and τ .

It is now easy to check that $F^{v \rightarrow G}$ is the graph of agreements of the permutation $\pi \in$

S_{n+m-1} defined by

$$\pi(i) \stackrel{\text{def}}{=} \begin{cases} \sigma(i), & \text{if } i < v \text{ and } \sigma(i) < \sigma(v), \\ \sigma(i) + m - 1, & \text{if } i < v \text{ and } \sigma(v) < \sigma(i), \\ \sigma(v) + \tau(i - v + 1) - 1, & \text{if } v \leq i < v + m, \\ \sigma(i - m + 1), & \text{if } v + m \leq i \text{ and } \sigma(i - m + 1) < \sigma(v), \\ \sigma(i - m + 1) + m - 1, & \text{if } v + m \leq i \text{ and } \sigma(v) < \sigma(i - m + 1). \end{cases}$$

In fact, the above shows that $\mathcal{M}[T_{\text{Perm}}]$ is weakly closed under substitutions, so $\mathcal{M}[T_{\text{PermGraph}}]$ inherits this property.

Now, by Theorem 5.4, it is sufficient to show that $\mathcal{M}[T_{\text{PermGraph}}]$ is not primally almost finite.

Recall that the family $\{G_n \mid n \geq 6\}$ of Example 3.21 is a family of prime graphs that is not almost finite. We claim that for every even⁸ $n \geq 6$, the graph G_n is a graph of agreements of some permutation. Indeed, G_n is the graph of agreements of the permutation $\pi_n \in S_{n+4}$ (see Figure 7) given by

$$\pi_n(i) \stackrel{\text{def}}{=} \begin{cases} n + 3, & \text{if } i = 1, \\ n + 1, & \text{if } i = 2, \\ n - 1, & \text{if } i = 3, \\ n + 4, & \text{if } i = 4, \\ n - i + 3, & \text{if } 6 \leq i \leq n \text{ and } i \text{ is even,} \\ n - i + 7, & \text{if } 5 \leq i \leq n - 1 \text{ and } i \text{ is odd,} \\ 1, & \text{if } i = n + 1, \\ 6, & \text{if } i = n + 2, \\ 4, & \text{if } i = n + 3, \\ 2, & \text{if } i = n + 4. \end{cases}$$

For example, the values of π_{14} (in sequence) are

$$17, 15, 13, 18, 16, 11, 14, 9, 12, 7, 10, 5, 8, 3, 1, 6, 4, 2.$$

Thus, $\mathcal{M}[T_{\text{PermGraph}}]$ is not primally almost finite, hence $T_{\text{PermGraph}} \notin \text{WR}$ by Theorem 5.4. \square

We conclude this section with an example of a universal theory T of graphs that is in WR essentially because of failure of the product condition of Proposition 5.7.

Proposition 5.13. *Consider the sequence of graphs $G = (C_{n^2+5})_{n \in \mathbb{N}}$ and let $T \stackrel{\text{def}}{=} \text{Th}(\phi_G)$ be the theory of positive models of the recursive blow-up ϕ_G relative to G (see Definition 4.7). Then $T \in \text{WR}$.*

⁸It is also true for odd n , but we only need an infinite subfamily, so even n suffices.

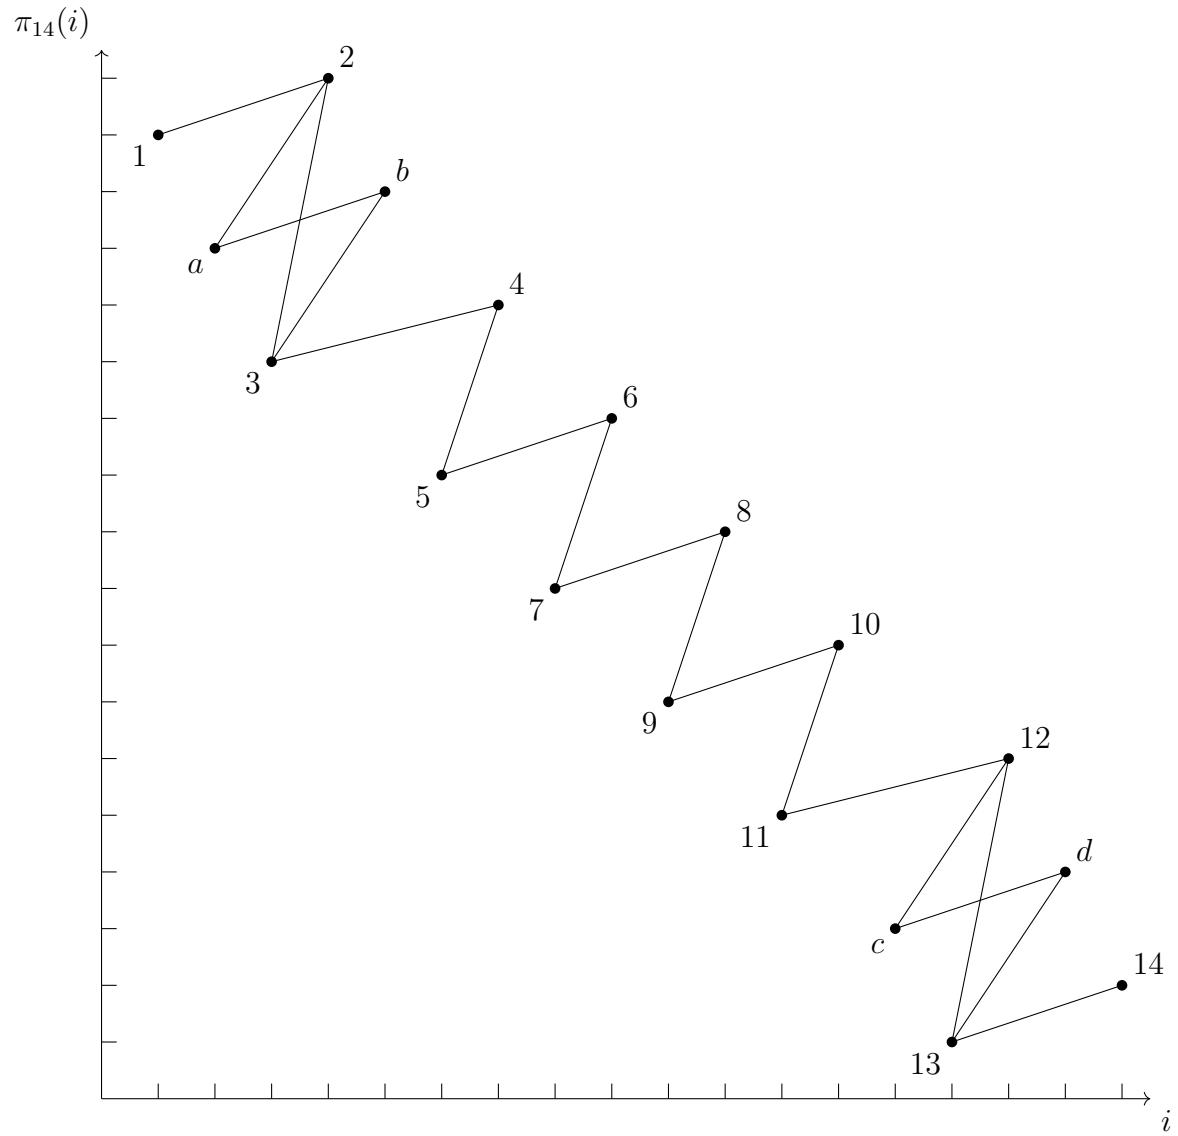


Figure 7: Graph of permutation π_{14} of proof of Proposition 5.12 represented as points (i.e., the set $\{(i, \pi_{14}(i)) \mid i \in [14]\}$). The edges of the corresponding graph of agreements G_{14} are represented as lines and the labels indicate the vertices of G_{14} .

Proof. Let $V = (V_n)_{n \in \mathbb{N}}$ be given by $V_n \stackrel{\text{def}}{=} V(C_{n^2+5})$ and let W^G be the graphon representation of ϕ_G given by (9). Since W^G is $\{0, 1\}$ -valued, we can view it as a continuum-sized graph H with vertex set Ω^V , consider the family \mathcal{F} of all finite graphs that are induced subgraphs of H .

We claim that $\mathcal{F} = Q(\phi_G)$. Indeed, we obviously have $Q(\phi_G) \subseteq \mathcal{F}$ and any induced subgraph $M \in \mathcal{F}$ of H must be an induced subgraph of the (finite) recursive blow-up $H'_{m_0} \stackrel{\text{def}}{=} R^{(C_5, C_6, C_9, \dots, C_{m_0^2+5})}$ for some $m_0 \in \mathbb{N}$ (see Definition 7.5) hence

$$\phi_G(M) \geq p(M, H'_{m_0}) \cdot \phi_G(H'_{m_0}) \geq p(M, H'_{m_0}) \cdot \frac{|H'_{m_0}|!}{|\text{Aut}(H'_{m_0})|} \cdot \left(\prod_{n=0}^{m_0} \frac{1}{n^2 + 5} \right)^{|H'_{m_0}|} > 0,$$

so $M \in Q(\phi_G)$.

Let us now show that $T \in \text{WR}$. Let $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ be an arbitrary limit of T . Since $\mathcal{M}[T] = \mathcal{F}$, there exists a sequence $U = (U_n)_{n \in \mathbb{N}}$ of finite subsets of Ω^V such that the sequence of finite graphs $(H|_{U_n})_{n \in \mathbb{N}}$ converges to ϕ .

For each $k \in \mathbb{N}$ and each $v \in V(C_{k^2+5})$, let

$$K_{k,v} \stackrel{\text{def}}{=} \{\sigma \in \Omega^V \mid \sigma_k = v\}.$$

For each $n \in \mathbb{N}$, let us construct a sequence $(U'_{n,k})_{k \in \mathbb{N}}$ of subsets of U_n inductively as follows. We set $U'_{n,0} \stackrel{\text{def}}{=} U_n$ and given $U'_{n,k}$, let $v_{n,k}$ be a vertex $v \in V(C_{k^2+5})$ that minimizes $|U'_{n,k} \cap K_{k,v}|$ (which can be zero) and let $U'_{n,k+1} \stackrel{\text{def}}{=} U'_{n,k} \setminus K_{n,v_{n,k}}$; note that

$$|U'_{n,k+1}| \geq \left(1 - \frac{1}{k^2 + 5}\right) \cdot |U'_{n,k}|.$$

We also let $U'_n \stackrel{\text{def}}{=} \bigcap_{k \in \mathbb{N}} U'_{n,k}$ and note that a simple induction gives

$$\frac{|U'_n|}{|U_n|} \geq \prod_{k \in \mathbb{N}} \left(1 - \frac{1}{k^2 + 5}\right) > 0. \quad (11)$$

Note also that the definition of U'_n implies that for every $k \in \mathbb{N}$, there exists $v \in V(C_{k^2+5})$ such that $U'_n \cap K_{k,v} = \emptyset$, which along with the definition of H implies that $H|_{U'_n} \in S(\{P_n \mid n \in \mathbb{N}\})$, where P_n is the path on n vertices.

Let then $(H|_{U'_{n_\ell}})_{\ell \in \mathbb{N}}$ be a convergent subsequence of $(H|_{U'_n})_{n \in \mathbb{N}}$ such that $(|U'_{n_\ell}|/|U_{n_\ell}|)_{\ell \in \mathbb{N}}$ is also convergent and let $\psi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ be the limit of $(H|_{U'_{n_\ell}})_{\ell \in \mathbb{N}}$. Then (11) implies that ψ is a sub-object of ϕ of measure at least $\prod_{k \in \mathbb{N}} (1 - 1/(k^2 + 5)) > 0$. But since $H|_{U'_{n_\ell}} \in S(\{P_n \mid n \in \mathbb{N}\})$, it follows that $\text{Th}(\psi)$ is primally almost finite, which by Lemma 5.6 gives $\text{Th}(\psi) \in \text{WR}$, so ψ has a weakly random sub-object, hence so does ϕ . \square

6 VC dimension and weak randomness

In this section we study how weak randomness and the class WR interact with the notion of VC dimension. We remind the reader that in this section we drop the qualifiers “weakly” and “strongly” from “closed under substitutions” as they are superfluous for graphs (see Remark 3.12).

Recall that for a non-trivial graph G , the *Vapnik–Chervonenkis dimension* [VC71] (VC dimension) of (neighborhoods of) G is the largest size $\text{VC}(G)$ of a set $U \subseteq V(G)$ that is *shattered* by neighborhoods of vertices of G in the sense that for every $A \subseteq U$, there exists $v \in V(G)$ such that $N_G(v) \cap U = A$, where $N_G(v) \stackrel{\text{def}}{=} \{w \in V(G) \mid G \models E(v, w)\}$ is the neighborhood of v in G . By convention, we also let $\text{VC}(K_0) \stackrel{\text{def}}{=} 0$.

Recall also that (the edge relation in) a class of graphs \mathcal{F} (or the corresponding universal theory $\text{Th}(\mathcal{F})$) is said to have *bounded VC dimension* if $\sup\{\text{VC}(G) \mid G \in \mathcal{F}\} < \infty$. In model theoretic language, the class has NIP (standing for *not the independence property*).

Finally, recall that by [LS10], a universal theory T of graphs has bounded VC dimension if and only if all graphons of T are a.e. $\{0, 1\}$ -valued⁹. Thus, studying VC dimension is directly related to studying whether the theory has fractional-valued graphons.

We start with a simple application of the theory of graph persistence developed so far.

Proposition 6.1. *If W is a graphon such that there exists a finite graph G with $\phi_W(G) = 0$, then W has an a.e. $\{0, 1\}$ -valued subgraphon W' . Furthermore, W' can be taken of the form $W' = W|_A$ for some positive measure set A .*

Proof. Define a $\{0, 1\}$ -valued graphon \widetilde{W} on the same space as W by

$$\widetilde{W}(x, y) \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if } 0 < W(x, y) < 1, \\ 0, & \text{if } W(x, y) \in \{0, 1\}. \end{cases}$$

Note that a subgraphon W' of W represented as $W|_f$ is a.e. $\{0, 1\}$ -valued if and only if $K_2 \notin Q(\widetilde{W}|_f)$. Hence, W has an a.e. $\{0, 1\}$ -valued subgraphon if and only if $K_2 \notin P(\widetilde{W})$. Thus, to prove the proposition, it is sufficient to show that $K_2 \in P(\widetilde{W})$ implies that every finite graph G has positive density in W .

Since $P(\widetilde{W})$ is closed under substitutions and induced subgraphs (Lemma 4.6), it follows that $K_n \in P(\widetilde{W})$ for every $n \in \mathbb{N}$ (as $K_n \cong K_2^{\rightarrow K_{n-1}}$). But note that each copy of K_n in \widetilde{W} corresponds to points of W whose pairs all have values in $(0, 1)$, hence have strictly fractional (conditional) probability of yielding edges, thus the fact that $K_{|G|}$ has positive density in \widetilde{W} implies that G has positive density in W .

For the final part, if $W' = W|_f$ for some function $f: X \rightarrow [0, 1]$, then letting $A \stackrel{\text{def}}{=} \{x \in X \mid f(x) > 0\}$ yields that $W|_A$ is an a.e. $\{0, 1\}$ -valued subgraphon of W . \square

⁹Let us warn the unfamiliarized reader that even if W is a.e. $\{0, 1\}$ -valued, its theory of positive graphs $\text{Th}(W)$ is not necessarily of bounded VC dimension. For example, the construction in the proof of Theorem 4.4 always yields a $\{0, 1\}$ -valued graphon with $\text{Th}(W) = \mathcal{F}$, even if $\mathcal{F} = \mathcal{M}[T_{\text{Graph}}]$, which clearly has unbounded VC dimension.

The reader might have noticed that with the notable exception of quasirandom graphons, all examples of weakly random graphons of Section 5 are $\{0, 1\}$ -valued. The next proposition says this is not a coincidence: every universal weakly random limit of a *proper* (strongly) persistent class of graphs \mathcal{F} (i.e., $\mathcal{F} \subsetneq \mathcal{M}[T_{\text{Graph}}]$) must necessarily be $\{0, 1\}$ -valued.

Theorem 6.2. *If W is a weakly random graphon such that there exists a finite graph G with $\phi_W(G) = 0$, then W is a.e. $\{0, 1\}$ -valued.*

Proof. We prove this by the contra-positive. Suppose W is a weakly random graphon over some space $\Omega = (X, \mathcal{A}, \mu)$ that is not a.e. $\{0, 1\}$ -valued. Let us show by induction in $|G|$ that every finite graph G has $\phi_W(G) > 0$. By possibly applying the Graphon Removal Lemma [Pet13, Theorem 1], it is enough to show that $T_{\text{ind}}(G, W) \not\subseteq \mathcal{D}_{V(G)}$.

Obviously $\phi_W(K_0) = \phi_W(K_1) = 1$. So suppose $|G| \geq 2$, let $v_0 \in V(G)$ and $H \stackrel{\text{def}}{=} G - v_0$. Since W is not a.e. $\{0, 1\}$ -valued, there exists $x_{v_0} \in [0, 1]$ such that the set

$$A \stackrel{\text{def}}{=} \{y \in [0, 1] \setminus \{x_{v_0}\} \mid 0 < W(x_{v_0}, y) < 1\}$$

has positive measure. Since $W|_A$ is a subgraphon of W , $W|_A$ is also weakly random and since by induction hypothesis, $\phi_W(H) > 0$ we get that there exists a point $(z, w) \in A^{V(H)} \times [0, 1]^{\binom{V(H)}{2}}$ that induces an off-diagonal copy of H in $W|_A$ (i.e., we have $(z, w) \in T_{\text{ind}}(H, W|_A)$).

Let us extend (z, w) to a point in $X^{V(G)} \times [0, 1]^{\binom{V(G)}{2}}$ by defining $z_{v_0} \stackrel{\text{def}}{=} x_{v_0}$ and

$$w_{\{v_0, w\}} \stackrel{\text{def}}{=} \begin{cases} 0, & \text{if } \{v_0, u\} \in E(G), \\ \frac{1 + W(x_{v_0}, z_u)}{2}, & \text{if } \{v_0, u\} \notin E(G), \end{cases}$$

for every $u \in V(H)$. It is straightforward to check that (z, w) yields an off-diagonal copy of G in W , concluding the proof. \square

Our next objective is to show that for families of graphs \mathcal{F} that are closed under substitutions and induced subgraphs, determining whether \mathcal{F} has bounded VC dimension is reduced to determining whether the family of prime graphs of \mathcal{F} has bounded VC dimension. To do so, we need a variation of the definition of VC dimension.

Definition 6.3. Given a non-trivial graph G , the *VC' dimension* of G (denoted $\text{VC}'(G)$) is the largest size of a set $U \subseteq V(G)$ that is *almost shattered* by the edge relation of G in the sense that for every non-empty $A \subsetneq U$, there exists $v \in V(G)$ such that $N_G(v) \cap U = A \setminus \{v\}$. By convention, we also let $\text{VC}'(K_0) \stackrel{\text{def}}{=} 0$.

Note that the notion of almost shattering is weaker than the notion of shattering in *two* points: we only care about sets A that are *non-empty proper subsets* of U and $N_G(v) \cap U$ only needs to match A up to possibly removing v from A . Note that for a non-trivial graph G , we always have $\text{VC}'(G) \geq 1$ as any singleton set is almost shattered by the edge relation of G .

Lemma 6.4. *For a non-trivial graph G , we have*

$$\max\{F(n) \mid n \leq \text{VC}'(G)\} \leq \text{VC}(G) \leq \text{VC}'(G),$$

where

$$F(n) \stackrel{\text{def}}{=} \begin{cases} \max \left\{ t \in \mathbb{N} \mid \sum_{i=0}^{t-1} \binom{n}{i} < \frac{2^n - 2}{n} \right\}, & \text{if } n \geq 3, \\ 0, & \text{if } n \leq 2. \end{cases}$$

Proof. Since the notion of shattering implies the notion of almost shattering, it follows trivially that $\text{VC}(G) \leq \text{VC}'(G)$.

Let now $n \leq \text{VC}'(G)$ and let us show that $F(n) \leq \text{VC}(G)$. Note that the result is trivial if $n \leq 2$ as $F(0) = F(1) = F(2) = 0$, so let us suppose $n \geq 3$. Since $n \leq \text{VC}'(G)$, we know that there exists a set $U \subseteq V(G)$ of size n that is almost shattered by the edge relation of G . For each non-empty $A \subsetneq U$, let $v_A \in V(G)$ be such that $N_G(v_A) \cap U = A \setminus \{v_A\}$ and let $\mathcal{F} \stackrel{\text{def}}{=} \{N_G(v_A) \cap U \mid \emptyset \neq A \subsetneq U\}$.

We claim that for every $B \subseteq U$, there are at most n non-empty sets $A \subsetneq U$ such that $N_G(v_A) \cap U = B$. Indeed, since $N_G(v_A) \cap U = A \setminus \{v_A\}$, the set of non-empty $A \subsetneq U$ with $N_G(v_A) \cap U = B$ must be contained in

$$\{B\} \cup \{B \cup \{u\} \mid u \in U\}.$$

When B is non-empty, the set above has size at most $|U| = n$ and when B is empty, the set above has size $n + 1$ but A cannot be equal to B .

Since there are $2^n - 2$ non-empty sets $A \subsetneq U$, we get $|\mathcal{F}| \geq (2^n - 2)/n$. On the other hand, by the definition of $F(n)$, we have

$$|\mathcal{F}| \geq \frac{2^n - 2}{n} > \sum_{i=0}^{F(n)-1} \binom{n}{i},$$

so by the Sauer–Shelah Lemma [Sau72, She72], the family \mathcal{F} shatters some $U' \subseteq U$ with $|U'| \geq F(n)$, thus $\text{VC}(G) \geq F(n)$. \square

Remark 6.5. It is easy to see that the function F of Lemma 6.4 is unbounded. Indeed, if there was a bound $t_0 \in \mathbb{N}$ such that $F(n) \leq t_0$ for every $n \in \mathbb{N}$, then we would have

$$\frac{2^n - 2}{n} \leq \sum_{i=0}^{t_0} \binom{n}{i} \leq (t_0 + 1) \cdot n^{t_0}$$

for every $n \in \mathbb{N}$, which yields a contradiction when n is sufficiently large.

As a corollary of Lemma 6.4, it then follows that a universal theory graphs T has bounded VC dimension if and only if there exists $k \in \mathbb{N}$ such that $\text{VC}'(G) \leq k$ for every graph G of T .

The next lemma shows that VC' dimension behaves very well with respect to the substitution operation.

Lemma 6.6. *Let F_1 and F_2 be non-trivial finite graphs and $v_0 \in V(F_1)$. Then we have $\text{VC}'(F_1^{v_0 \rightarrow F_2}) = \max\{\text{VC}'(F_1), \text{VC}'(F_2)\}$.*

Proof. Without loss of generality, suppose $V(F_1) \cap V(F_2) = \emptyset$ and let $G \stackrel{\text{def}}{=} F_1^{v_0 \rightarrow F_2}$. Since both F_1 and F_2 are induced subgraphs of G (as both F_1 and F_2 are non-trivial), it follows that $\text{VC}'(G) \geq \max\{\text{VC}'(F_1), \text{VC}'(F_2)\}$.

To prove the other inequality, let $U \subseteq V(G)$ be a set that is almost shattered by the edge relation of G with $|U| = \text{VC}'(G)$.

Suppose first that $U \subseteq V(F_i)$ for some $i \in [2]$. Then we claim that the edge relation of F_i also almost shatters U . Indeed, if $A \subsetneq U$ is a non-empty set, then we know that there exists $u \in V(G)$ such that $N_G(u) \cap U = A \setminus \{u\}$. Since $A \neq \emptyset$ and $A \neq V(F_i)$ (as $A \subsetneq U \subseteq V(F_i)$), we must have $u \in V(F_i)$ (as every $u \in V(F_{3-i})$ is either adjacent to all of $V(F_i)$ or not adjacent to all of $V(F_i)$) and thus $N_{F_i}(u) \cap U = A \setminus \{u\}$. Therefore, in this case, we get $\text{VC}'(G) \leq \text{VC}'(F_i) \leq \max\{\text{VC}'(F_1), \text{VC}'(F_2)\}$.

Suppose then that $U \not\subseteq V(F_1)$ and $U \not\subseteq V(F_2)$. Then we claim that $|U \cap V(F_2)| = 1$. Suppose not and let $v_1 \in U \cap V(F_1)$ and $v_2, w_2 \in U \cap V(F_2)$ with $v_2 \neq w_2$. We consider first the case when $\{v_0, v_1\} \in E(F_1)$ (recall that v_0 is the vertex of F_1 that is being substituted: $G = F_1^{v_0 \rightarrow F_2}$), let $A \stackrel{\text{def}}{=} \{v_2\} \subsetneq U$ and let $u \in V(G)$ be such that $N_G(u) \cap U = A \setminus \{u\}$. Since $\{v_0, v_1\} \in E(F_1)$ and $v_1 \notin A$, we must have $u \notin V(F_2)$ (as every vertex of $V(F_2)$ is adjacent to v_1 in G) and since $w_2 \notin A$, we must have $u \notin V(F_1)$ (as every vertex of $V(F_1)$ is adjacent to w_2 in G), a contradiction. Consider then the case when $\{v_0, v_1\} \notin E(F_1)$, let $A \stackrel{\text{def}}{=} \{v_1, v_2\} \subsetneq U$ and let $u \in V(G)$ be such that $N_G(u) \cap U = A \setminus \{u\}$. Since $\{v_0, v_1\} \notin E(F_1)$ and $v_1 \in A$, we must have $u \notin V(F_2)$ and since $v_2 \in A$, we must have $u \notin V(F_1)$, a contradiction. This concludes the proof of the claim, that is, we have $|U \cap V(F_2)| = 1$.

Let then w_0 be the unique element of $U \cap V(F_2)$. We now claim that the set $U' \stackrel{\text{def}}{=} (U \setminus \{w_0\}) \cup \{v_0\}$ is almost shattered by the edge relation of F_1 . Let $A' \subsetneq U'$ be a non-empty set and let

$$A \stackrel{\text{def}}{=} \begin{cases} A', & \text{if } v_0 \notin A', \\ (A' \setminus \{v_0\}) \cup \{w_0\}, & \text{if } v_0 \in A'. \end{cases}$$

Then there exists $u \in V(G)$ such that $N_G(u) \cap U = A \setminus \{u\}$. Consider first the case when $u \in V(F_1)$. Then we have

$$N_{F_1}(u) \cap U' = \begin{cases} N_G(u) \cap U, & \text{if } w_0 \notin N_G(u) \cap U, \\ ((N_G(u) \cap U) \setminus \{w_0\}) \cup \{v_0\}, & \text{if } w_0 \in N_G(u) \cap U, \end{cases}$$

hence $N_{F_1}(u) \cap U' = A' \setminus \{u\}$ (as $N_G(u) \cap U = A \setminus \{u\}$). Consider now the case when $u \notin V(F_1)$ and note that

$$N_{F_1}(v_0) \cap U' = (N_G(u) \cap U) \setminus \{w_0\} = A \setminus \{u, w_0\} = A' \setminus \{v_0\}$$

since $u \notin A'$ (as $A' \subseteq V(F_1)$). Thus U' is almost shattered by the edge relation of F_1 , hence $\text{VC}'(G) \leq \text{VC}'(F_1) \leq \max\{\text{VC}'(F_1), \text{VC}'(F_2)\}$. \square

The following simple consequence of Lemmas 6.4 and 6.6 (and Remark 6.5) says that when $\mathcal{F} = S(\mathcal{F}')$, then \mathcal{F} has bounded VC dimension if and only if \mathcal{F}' has bounded VC dimension.

Theorem 6.7. *Let \mathcal{F} and \mathcal{F}' be families of finite graphs up to isomorphism and suppose $\mathcal{F} = S(\mathcal{F}')$. Then \mathcal{F} has bounded VC dimension if and only if \mathcal{F}' has bounded VC dimension.*

In particular, if \mathcal{F} is a family of finite graphs that is closed under substitutions and under induced subgraphs and \mathcal{P} is the family of all prime graphs of \mathcal{F} , then \mathcal{F} has bounded VC dimension if and only if \mathcal{P} has bounded VC dimension.

Proof. By Lemmas 6.4 and 6.6 and Remark 6.5 with a simple induction, we have

$$\begin{aligned} \sup\{\text{VC}(F) \mid F \in \mathcal{F}\} < \infty &\iff \sup\{\text{VC}'(F) \mid F \in \mathcal{F}\} < \infty \\ &\iff \sup\{\text{VC}'(F') \mid F' \in \mathcal{F}'\} < \infty \\ &\iff \sup\{\text{VC}(F') \mid F' \in \mathcal{F}'\} < \infty, \end{aligned}$$

so the first statement follows.

The second statement follows from the first one along with Lemma 3.5. \square

As a direct corollary of Theorem 6.7, it follows that any primally finite family \mathcal{F} that is closed under substitutions and under induced subgraphs has bounded VC dimension. Our next objective is to show that the same is true in the primally almost finite case. Before we do so, we need yet another example of a family of prime graphs that is not almost finite.

Example 6.8. For each $n \geq 9$ odd, let G'_n be the graph obtained from the path on n vertices P_n by adding two vertices a and b adjacent precisely to the fourth and fourth from last vertices of P_n , respectively and connecting all even vertices into a clique (see Figure 8). Formally, we have

$$\begin{aligned} V(G'_n) &\stackrel{\text{def}}{=} [n] \cup \{a, b\}, \\ E(G'_n) &\stackrel{\text{def}}{=} \{\{i, i+1\} \mid i \in [n-1]\} \cup \{a, 4\} \cup \{b, n-3\} \\ &\quad \cup \{\{2i, 2j\} \mid i, j \in [(n-1)/2] \wedge i \neq j\}. \end{aligned}$$

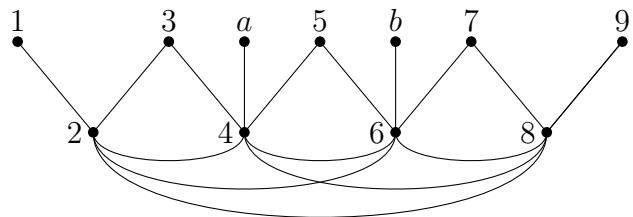
It is straightforward to check that $\{G'_n \mid n \geq 9 \text{ odd}\}$ is a family of prime graphs that is not almost finite.

Before we prove the theorem, we need a small consequence of Ramsey's Theorem.

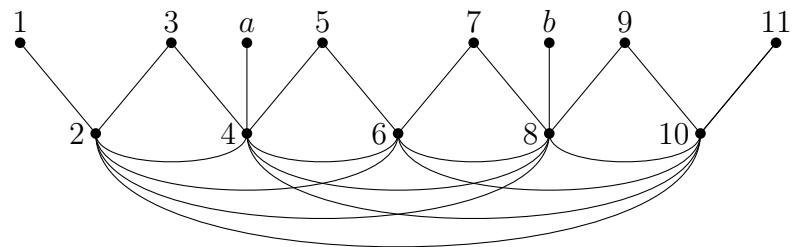
Lemma 6.9. *For every $n \in \mathbb{N}$ and every graphon W , we have*

$$\phi_W(K_n) + \phi_W(\overline{K}_n) \geq \binom{R(n, n)}{n}^{-1},$$

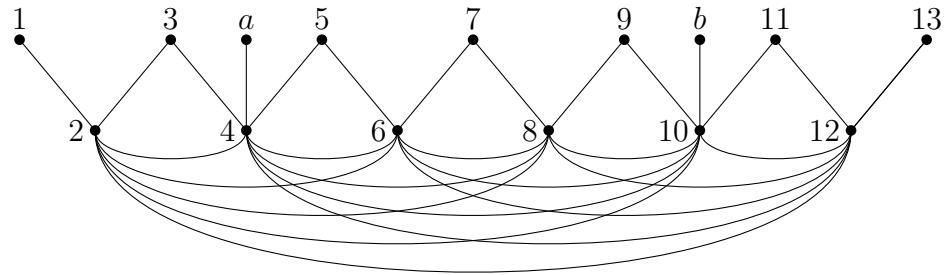
where $R(n, n)$ is the (n, n) -Ramsey number.



$G'_9.$



$G'_ {11}.$



$G'_ {13}.$

Figure 8: Prime graphs G'_n of Example 6.8 that form a family that is not almost finite.

Proof. We have

$$\begin{aligned}\phi_W(K_n) + \phi_W(\overline{K}_n) &= \sum_{M \in \mathcal{M}_{R(n,n)}[T_{\text{Graph}}]} (p(K_n, M) + p(\overline{K}_n, M)) \cdot \phi_W(M) \\ &\geq \binom{R(n,n)}{n}^{-1} \sum_{M \in \mathcal{M}_{R(n,n)}[T_{\text{Graph}}]} \phi_W(M) = \binom{R(n,n)}{n}^{-1},\end{aligned}$$

where the inequality follows since at least one n -sized subset of each $M \in \mathcal{M}_{R(n,n)}[T_{\text{Graph}}]$ must induce either K_n or \overline{K}_n in M . \square

Theorem 6.10. *Let T be a universal theory of graphs. If $\mathcal{M}[T]$ is primally almost finite, then the edge relation in T has NIP (i.e., bounded VC dimension).*

In particular, the edge relation in every universal theory of graphs $T' \in \text{WR}$ such that $\mathcal{M}[T']$ is closed under substitutions has NIP.

Proof. The second assertion follows from the first along with Theorem 5.4.

We prove the first assertion by the contra-positive. Since the edge relation in T has unbounded VC dimension, by [LS10], there exists a graphon W that is a limit of T and is *not* a.e. $\{0, 1\}$ -valued. By possibly applying the Graphon Removal Lemma [Pet13, Theorem 1], we may suppose that every graph G that has an off-diagonal copy in W has positive density in W .

Our objective is to present a family of prime graphs that is not almost finite and such that all graphs in this family have an off-diagonal copy in W (thus $\mathcal{M}[T]$ is not primally almost finite).

For this purpose, we will show that for each $n \in \mathbb{N}$ with $n \geq 5$, one of the following graphs appears as an off-diagonal copy in W :

- i. The graph G_{2n-4} of Example 3.21.
- ii. The complement \overline{G}_{2n-4} of the graph of Example 3.21.
- iii. The graph G'_{2n-1} of Example 6.8.

Since each of these families is a family of prime graphs that is not almost finite (note that primality is preserved under complementation) and one of them must occur for infinitely many n , it will follow that $\mathcal{M}[T]$ is not primally almost finite as desired.

Without loss of generality, let us suppose that the underlying space of W is $[0, 1]$ and let $(x_0, y_0) \in (0, 1)^2$ be a Lebesgue density point with respect to ℓ^∞ -balls of the positive measure set $A \stackrel{\text{def}}{=} W^{-1}((0, 1))$ with $x_0 \neq y_0$.

Let $\epsilon > 0$ be such that $\epsilon < (n \cdot \binom{R(n,n)}{n})^{-2}$ and let $\delta > 0$ be small enough so that

$$\frac{\lambda(A \cap B_\delta(x_0, y_0))}{\lambda(B_\delta(x_0, y_0))} \geq 1 - \epsilon,$$

where $B_\delta(x_0, y_0)$ is the ℓ^∞ -ball of radius δ centered at (x_0, y_0) . We may also suppose that $\delta > 0$ is small enough so that $(x_0 - \delta, x_0 + \delta)$ and $(y_0 - \delta, y_0 + \delta)$ are disjoint subsets of $[0, 1]$.

Consider the set

$$C \stackrel{\text{def}}{=} \{(x, y) \in B_\delta(x_0)^n \times B_\delta(y_0)^n \mid \forall i, j \in [n], (x_i, y_j) \in A\}.$$

With a simple union bound, we have

$$\lambda(C) \geq (1 - n^2\epsilon) \cdot \lambda(B_\delta(x_0, y_0))^n > \left(1 - \binom{R(n, n)}{n}^{-2}\right) \lambda(B_\delta(x_0, y_0))^n. \quad (12)$$

Define also

$$\begin{aligned} C' \stackrel{\text{def}}{=} & \{(x, y) \in B_\delta(x_0)^n \times B_\delta(y_0)^n \mid \exists z \in [0, 1]^{\binom{[n]}{2}}, (x, z) \in T_{\text{ind}}(K_n, W) \cup T_{\text{ind}}(\overline{K}_n, W) \\ & \quad \wedge \exists w \in [0, 1]^{\binom{[n]}{2}}, (y, w) \in T_{\text{ind}}(K_n, W) \cup T_{\text{ind}}(\overline{K}_n, W)\}. \end{aligned}$$

By Lemma 6.9, we have

$$\lambda(C') \geq \binom{R(n, n)}{n}^{-2} \cdot \lambda(B_\delta(x_0, y_0))^n.$$

Putting this together with (12), we conclude that $\lambda(C \cap C') > 0$.

Let then $(x, y) \in C \cap C'$ be a point with all coordinates distinct. We now consider four cases.

Case 1. There exist points $z, w \in [0, 1]^{\binom{[n]}{2}}$ such that $(x, z), (y, w) \in T_{\text{ind}}(\overline{K}_n, W)$. In this case, we construct an off-diagonal copy $(\hat{x}, \hat{y}) \in [0, 1]^{V(G_{2n-4})} \times [0, 1]^{\binom{V(G_{2n-4})}{2}}$ of the graph G_{2n-4} of Example 3.21 as follows. Recall that $V(G_{2n-4}) = [2n-4] \cup \{a, b, c, d\}$ and for convenience of notation, let us make the identifications $a \stackrel{\text{def}}{=} 2n-3, b \stackrel{\text{def}}{=} 2n-2, c \stackrel{\text{def}}{=} 2n-1$ and $d \stackrel{\text{def}}{=} 2n$. Then for each $i, j \in V(G_{2n-4})$ with $i \neq j$, let

$$\begin{aligned} \hat{x}_i &\stackrel{\text{def}}{=} \begin{cases} x_{i/2}, & \text{if } i \in [2n] \text{ is even,} \\ y_{(i+1)/2}, & \text{if } i \in [2n] \text{ is odd,} \end{cases} \\ \hat{y}_{\{i, j\}} &\stackrel{\text{def}}{=} \begin{cases} z_{\{i/2, j/2\}}, & \text{if } i, j \in [2n] \text{ are both even,} \\ w_{\{(i+1)/2, (j+1)/2\}}, & \text{if } i, j \in [2n] \text{ are both odd,} \\ 0, & \text{if } i \in [2n] \text{ is even, } j \in [2n] \text{ is odd and } \{i, j\} \in E(G_{2n-4}), \\ \frac{1 + W(x_{i/2}, y_{(j+1)/2})}{2}, & \text{if } i \in [2n] \text{ is even, } j \in [2n] \text{ is odd and } \{i, j\} \notin E(G_{2n-4}). \end{cases} \end{aligned}$$

The fact that $(x, y) \in C \cap C'$ and all coordinates of (x, y) are distinct guarantees that (\hat{x}, \hat{y}) is an off-diagonal copy of G_{2n-4} .

Case 2. There exist points $z, w \in [0, 1]^{n \choose 2}$ such that $(x, z), (y, w) \in T_{\text{ind}}(K_n, W)$. In this case, a construction analogous to the one in case 1 yields an off-diagonal copy of the complement \overline{G}_{2n-4} of the graph of Example 3.21.

Case 3. There exist points $z, w \in [0, 1]^{n \choose 2}$ such that $(x, z) \in T_{\text{ind}}(K_n, W)$ and $(y, w) \in T_{\text{ind}}(\overline{K}_n, W)$. We construct an off-diagonal copy $(\widehat{x}, \widehat{y}) \in [0, 1]^{V(G'_{2n-1})} \times [0, 1]^{V(G'_{2n-1})}$ of the graph G'_{2n-1} of Example 6.8 as follows. Recall that $V(G'_{2n-1}) = [2n-1] \cup \{a, b\}$ and for convenience of notation, let us make the identifications $a \stackrel{\text{def}}{=} 2n+1$ and $b \stackrel{\text{def}}{=} 2n+3$ (nothing gets identified with the points $2n$ and $2n+2$). Then for each $i, j \in V(G'_{2n-1})$ with $i \neq j$, let

$$\widehat{x}_i \stackrel{\text{def}}{=} \begin{cases} x_{i/2}, & \text{if } i \in [2n-1] \text{ is even,} \\ y_{(i+1)/2}, & \text{if } i \in [2n+3] \text{ is odd,} \end{cases}$$

$$\widehat{y}_{\{i,j\}} \stackrel{\text{def}}{=} \begin{cases} z_{\{i/2, j/2\}}, & \text{if } i, j \in [2n-1] \text{ are both even,} \\ w_{\{(i+1)/2, (j+1)/2\}}, & \text{if } i, j \in [2n+3] \text{ are both odd,} \\ 0, & \text{if } i \in [2n-1] \text{ is even, } j \in [2n+3] \text{ is odd and } \{i, j\} \in E(G'_{2n-1}), \\ \frac{1 + W(x_{i/2}, y_{(j+1)/2})}{2}, & \text{if } i \in [2n-1] \text{ is even, } j \in [2n+3] \text{ is odd and } \{i, j\} \notin E(G'_{2n-1}). \end{cases}$$

The fact that $(x, y) \in C \cap C'$ and all coordinates of (x, y) are distinct guarantees that $(\widehat{x}, \widehat{y})$ is an off-diagonal copy of G'_{2n-1} .

Case 4. There exist points $z, w \in [0, 1]^{n \choose 2}$ such that $(x, z) \in T_{\text{ind}}(\overline{K}_n, W)$ and $(y, w) \in T_{\text{ind}}(K_n, W)$. This case follows from case 3 by swapping the roles of x and y .

Therefore $\mathcal{M}[T]$ is not primally almost finite. \square

Remark 6.11. The assumption that $\mathcal{M}[T]$ is closed under substitution is crucial for the second part of Theorem 6.10, since, for example, the universal theory $T_{\text{Bipartite}}$ of bipartite graphs clearly is in $\text{AEHP} \subseteq \text{WR}$ (as every limit contains an empty subgraphon of measure at least $1/2$) but has unbounded VC dimension.

For an example of a theory with unbounded VC dimension that is in $\text{WR} \setminus \text{AEHP}$, let T_{C_4} be the universal theory of graphs that are induced subgraphs of some (finite) recursive blow-up of C_4 (which has bounded VC dimension by Theorem 6.10 as T_{C_4} is primally finite), let \mathcal{F} be the family of graphs G such that there exists a partition $V(G) = A \cup B$ such that both $G|_A$ and $G|_B$ are models of T_{C_4} and let $T \stackrel{\text{def}}{=} \text{Th}(\mathcal{F})$ be the corresponding universal theory of graphs. Obviously, every bipartite graph is a model of T , so T has unbounded VC dimension.

Since at least one of A or B must have at least half of the vertices, it follows that every limit of T has a subgraphon that is a limit of T_{C_4} and since $T_{C_4} \in \text{WR}$ (by Proposition 5.9), we get $T \in \text{WR}$. On the other hand, since every model of T_{C_4} is a model of T and $T_{C_4} \notin \text{AEHP}$ (by [CM22, Lemma 8.8] or Proposition 5.9 again), it follows that $T \notin \text{AEHP}$.

7 Persistence for universal theories

In this section, we generalize the results of Section 4 on (strongly) persistent classes to arbitrary universal theories in finite relational languages. Table 1 below contains the correspondence between the theorems and lemmas of Sections 4 and 5 and their generalizations in this and the next section.

Graph result	Universal theory result	Drawbacks
Lemma 4.3	Lemma 7.2(i)	None.
Theorem 4.4	Theorem 7.9	Persistence (item (iii)) can only be added to the list of equivalences when all arities are at most 2.
Lemma 4.6	Lemma 7.3	Requires either that all arities are at most 2 (item (i)) or weak randomness (item (ii)).
Lemma 4.12	Lemma 7.6 and Proposition 7.8(i),(ii)	When the recursive blow-ups are not conservative (see Definitions 7.5 and 7.7), only partial information is known about the limit theor.
Lemma 4.13	Proposition 7.8(iii)	None.
Lemma 5.3	Lemma 7.2(ii)	None.
Theorem 5.4	Propositions 8.2 and 8.4	Backward direction requires all arities to be at most 2. Forward direction is trivial if all arities are at least 3.
Lemma 5.6	Proposition 8.2	Requires all arities to be at most 2.
Proposition 5.7	Propositions 7.8(iv) and 8.4	When all arities are at least 3, there are only finitely many prime structures (see Remark 3.13).

Table 1: Correspondence between theorems and lemmas of Sections 4 and 5 and their generalizations in Sections 7 and 8. Some generalizations have drawbacks (e.g., extra hypotheses or the result might be trivial) as pointed out in the third column.

Definition 7.1. Let $\phi \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$. The set of *positive \mathcal{L} -structures in ϕ* is the set $Q(\phi) \stackrel{\text{def}}{=} \mathcal{M}[\text{Th}(\phi)]$ of all finite \mathcal{L} -structures M (up to isomorphism) such that $\phi(M) > 0$. The set of *persistently positive \mathcal{L} -structures in ϕ* is the set $P(\phi) \stackrel{\text{def}}{=} \bigcap_{\psi} Q(\psi)$, where the intersection is over all sub-objects of ϕ . We extend these definitions naturally to Euclidean structures \mathcal{N} in \mathcal{L} by $Q(\mathcal{N}) \stackrel{\text{def}}{=} Q(\phi_{\mathcal{N}})$ and $P(\mathcal{N}) \stackrel{\text{def}}{=} P(\phi_{\mathcal{N}})$.

We say that ϕ is *weakly random* if $P(\phi) = Q(\phi)$.

A family \mathcal{F} of finite \mathcal{L} -structures (up to isomorphism) is called *persistent* if there exists ϕ such that $P(\phi) = \mathcal{F}$. The family \mathcal{F} is called *strongly persistent* if there exists a weakly random ϕ such that $P(\phi) = \mathcal{F}$ (which must also equal $Q(\phi)$); in this case, we also say that ϕ is a *universal weakly random limit of \mathcal{F}* .

Lemma 7.2. *Let \mathcal{N} be an Euclidean structure in \mathcal{L} over $\Omega = (X, \mathcal{A}, \mu)$. Then the following hold.*

- i. $P(\mathcal{N}) = \bigcap_A Q(\mathcal{N}|_A^F)$, where the intersection is over all positive measure $A \in \mathcal{A}$ and all measure-isomorphisms F modulo 0 from Ω_A to Ω (equivalently, we can also use a single measure-isomorphism F_A modulo 0 for each positive measure $A \in \mathcal{A}$).
- ii. $\phi_{\mathcal{N}}$ has a weakly random sub-object if and only if there exists a positive measure $A \in \mathcal{A}$ and a measure-isomorphism F modulo 0 from Ω_A to Ω such that $\phi_{\mathcal{N}|_A^F}$ is weakly random.

Proof. Both items follow from the fact that if $f: X \rightarrow [0, 1]$ is a measurable function with $\int_X f \, d\mu > 0$ and F is a measure-isomorphism modulo 0 from Ω_f to Ω , then for $A = \{x \in X \mid f(x) > 0\}$ and any measure-isomorphism \tilde{F} modulo 0 from Ω_A to Ω , we have $Q(\mathcal{N}|_f^F) = Q(\mathcal{N}|_A^{\tilde{F}})$. \square

Lemma 7.3. *The following hold for $\phi \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$.*

- i. *If all predicate symbols of \mathcal{L} have arity at most 2, then $P(\phi)$ is strongly closed under substitutions and closed under substructures.*
- ii. *If ϕ is weakly random, then $P(\phi)$ is weakly closed under substitutions and closed under substructures.*

Proof. Since obviously $K_0 \in P(\phi)$, by Lemma 3.15, it is sufficient to show the assertions of closed under substitutions in each item.

Let $F_1, F_2 \in P(\phi)$ and $v \in V(F_1)$ and let \mathcal{F} be the set of standard substitutions of v in F_1 by F_2 .

Note that in item (i), by Remark 3.12, \mathcal{F} has a unique element (and the notion of strongly and weakly closed under substitutions coincide). Thus, in both items, our objective is to show that $\mathcal{F} \cap P(\phi)$ is non-empty.

Let ψ be a sub-object of ϕ and let \mathcal{N} be an Euclidean structure in \mathcal{L} over some space $\Omega = (X, \mathcal{A}, \mu)$ with $\phi_{\mathcal{N}} = \psi$. We claim that $\mathcal{F} \cap Q(\psi)$ is non-empty. Suppose not, that is, suppose $t_{\text{ind}}(F, \mathcal{N}) = 0$ for every $F \in \mathcal{F}$. By possibly applying the Induced Euclidean Removal Lemma [CR20, Theorem 3.3], we may suppose that $T_{\text{ind}}(F, \mathcal{N}) \subseteq \mathcal{D}_V$ for every $F \in \mathcal{F}$.

Since $F_1 \in P(\phi)$, we must have $F_1 \in Q(\mathcal{N})$, that is, we have $t_{\text{ind}}(F_1, \mathcal{N}) > 0$. For every $x \in X^{r(V(F_1)) \setminus \{v\}}$, let

$$U_x \stackrel{\text{def}}{=} \{y \in X \mid (x, y) \in T_{\text{ind}}(F_1, \mathcal{N})\}.$$

By Fubini's Theorem, there exists $x \in X^{r(V(F_1)) \setminus \{v\}}$ with all coordinates distinct such that $\mu(U_x) > 0$. Let G be a measure-isomorphism modulo 0 from Ω_{U_x} to Ω and since $\phi_{\mathcal{N}|_{U_x}^G}$ is a sub-object of ψ , hence also of ϕ , we must have $F_2 \in Q(\mathcal{N}|_{U_x}^G)$, which implies that there exists $z \in \mathcal{E}_{V(F_2)}$ such that

- a. all coordinates of z are distinct;
- b. all coordinates of z are distinct from the coordinates of x ;
- c. for every $v \in V(F_2)$, we have $z_{\{v\}} \in U_x$;
- d. we have $z \in T_{\text{ind}}(F_2, \mathcal{N})$.

Define then the point $w \in \mathcal{E}_V$ by the following procedure.

1. For each $A \subseteq r(V(F_1 - v))$, let $w_A \stackrel{\text{def}}{=} x_A$.
2. For each $A \subseteq r(V(F_1 - v))$ and each $u \in V(F_2)$, let $w_{A \cup \{u\}} \stackrel{\text{def}}{=} x_{A \cup \{u\}}$.
3. For each $A \subseteq r(V(F_2))$, let $w_A \stackrel{\text{def}}{=} z_A$.
4. Define all other coordinates of w arbitrarily.

Note that all coordinates of w that are indexed by single vertices get defined in items (1) and (3) and their definitions guarantee that they are distinct from each other, that is, we have $w \notin \mathcal{D}_V$. Let then F be the unique \mathcal{L} -structure with $w \in T_{\text{ind}}(F, \mathcal{N})$. Then items (1) and (2) ensure that all injections $V(F_1) \rightarrow V$ acting identically on $V(F_1 - v)$ are embeddings of F_1 in F and item (3) ensures that the injection $V(F_2) \rightarrow V$ that acts identically on $V(F_2)$ is an embedding of F_2 in F . Thus, we must have $F \in \mathcal{F}$.

Therefore, we have showed that for every sub-object ψ of ϕ , we have $\mathcal{F} \cap Q(\psi) \neq \emptyset$.

In item (i), since \mathcal{F} has a single element F , it follows that $F \in P(\phi)$, hence $P(\phi)$ is strongly (in this case, equivalently, weakly) closed under substitutions.

In item (ii), since $Q(\psi) = P(\phi)$ as ϕ is weakly random, it follows that $\mathcal{F} \cap P(\phi) \neq \emptyset$, so $P(\phi)$ is weakly closed under substitutions. \square

The next example shows why the hypotheses of Lemma 7.3 to get $P(\phi)$ weakly closed under substitutions are crucial.

Example 7.4. Consider $\phi \in \text{Hom}^+(\mathcal{A}[T_3\text{-Hypergraph}], \mathbb{R})$ that is the disjoint union of a clique and an anti-clique of the same size, that is, $\phi = \phi_{\mathcal{N}}$ for the $T_3\text{-Hypergraph}$ -on \mathcal{N} over $[0, 1]$ given by

$$\mathcal{N}_E \stackrel{\text{def}}{=} \left\{ x \in \mathcal{E}_3 \mid \max\{x_{\{1\}}, x_{\{2\}}, x_{\{3\}}\} < \frac{1}{2} \right\}.$$

Since ϕ contains both a clique and an anti-clique of positive measure, it follows that $P(\phi)$ does not contain any models of size at least 3. However, since $T_3\text{-Hypergraph}$ is 2-categorical, $P(\phi)$ must contain the unique model $K_2^{(3)}$ of size 2. It then follows that $P(\phi)$ is not even weakly closed under substitutions as any substitution of any vertex of $K_2^{(3)}$ by $K_2^{(3)}$ must have size 3.

Definition 7.5. Given a finite sequence $N = (N_0, \dots, N_n)$ of finite \mathcal{L} -structures with $|N_i| \geq 2$ for every $i \in \{0, \dots, n\}$, a *recursive blow-up relative to N* is an \mathcal{L} -structure R with $V(R) \stackrel{\text{def}}{=} \prod_{i=0}^n V(N_i)$ such that for every $j \in \{0, \dots, n\}$ and every $\sigma \in \prod_{i=0}^{j-1} V(N_i)$, every function $f: V(N_j) \rightarrow V(R)$ such that $f(v)|_{\{0, \dots, j-1\}} = \sigma$ and $f(v)_j = v$ for every $v \in V(N_j)$ is an embedding of N_j in R .

The unique recursive blow-up R relative to N that has the smallest possible relation sets P^R ($P \in \mathcal{L}$) is called the *conservative recursive blow-up relative to N* and is denoted R^N . Formally, it is given by $V(R^N) \stackrel{\text{def}}{=} \prod_{i=0}^n V(N_i)$ and

$$P^{R^N} \stackrel{\text{def}}{=} \left\{ (\sigma, \alpha_j, \tau^j)_{j=0}^{k(P)} \in (V(R^N))_{k(P)} \mid \sigma \in \prod_{\ell=0}^{i-1} V(N_\ell) \wedge \alpha \in P^{N_i} \wedge \forall j \in [k(P)], \tau^j \in \prod_{\ell=i+1}^n V(N_\ell) \right\}$$

for every $P \in \mathcal{L}$.

Given an infinite sequence $N = (N_i)_{i \in \mathbb{N}}$ of finite \mathcal{L} -structures with $|N_i| \geq 2$ for every $i \in \mathbb{N}$, a *compatible sequence of recursive blow-ups relative to N* is a sequence $R = (R_i)_{i \in \mathbb{N}}$ such that

- i. for every $i \in \mathbb{N}$, R_i is a recursive blow-up relative to (N_0, \dots, N_i) ;
- ii. for every $i \in \mathbb{N}$, every function $f: V(R_i) \rightarrow V(R_{i+1})$ such that $f(v)|_{\{0, \dots, i\}} = v$ is an embedding of R_i in R_{i+1} .

We call R *conservative* if further $R_i = R^{(N_0, \dots, N_i)}$ (it is easy to see that this is always compatible).

Lemma 7.6. Let \mathcal{F} be a family of finite \mathcal{L} -structures that is weakly closed under substitutions and closed under substructures and let $N = (N_i)_{i \in \mathbb{N}}$ be a sequence in \mathcal{F} with $|N_i| \geq 2$ for every $i \in \mathbb{N}$. Then there exists a compatible sequence $R = (R_i)_{i \in \mathbb{N}}$ of recursive blow-ups relative to N with $R_i \in \mathcal{F}$ for every $i \in \mathbb{N}$.

Proof. We construct the compatible sequence $R = (R_i)_{i \in \mathbb{N}}$ inductively by setting $R_0 \stackrel{\text{def}}{=} N_0$ and given R_i , we enumerate the vertices of R_i as $v_1^i, \dots, v_{t_i}^i$, inductively define $F_0^i, \dots, F_{t_i}^i$ by $F_0^i \stackrel{\text{def}}{=} R_i$, let $F_{j+1}^i \in \mathcal{F}$ be a standard substitution of v_{j+1}^i in F_j^i by N_{i+1} and set $R_{i+1} \stackrel{\text{def}}{=} F_{t_i}^i$. It is straightforward to check by induction that R is a compatible sequence of recursive blow-ups relative to N with $R_i \in \mathcal{F}$ for every $i \in \mathbb{N}$. \square

Definition 7.7. Given an infinite sequence $N = (N_i)_{i \in \mathbb{N}}$ of finite \mathcal{L} -structures with $|N_i| \geq 2$ for every $i \in \mathbb{N}$, we let the *conservative recursive blow-up relative to N* be the $T_{\mathcal{L}}$ -on \mathcal{N}^N defined as follows. We let $V = (V_\ell)_{\ell \in \mathbb{N}}$ be defined by $V_\ell \stackrel{\text{def}}{=} V(N_\ell)$ and we define \mathcal{N}^N over the Cantor probability space $\Omega^V = (\prod_{\ell \in \mathbb{N}} V_\ell, \mathcal{A}, \nu^V)$ (see Definition 4.7) by

$$\mathcal{N}_P^N \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{k(P)}(\Omega^V) \mid \exists i \in \mathbb{N}, R^{(N_0, \dots, N_i)} \models P(t_i^P(x))\},$$

where $t_i^P(x) \in (\prod_{\ell=0}^i V_\ell)^{k(P)}$ is given by

$$t_i^P(x)_j \stackrel{\text{def}}{=} x_{\{j\}}|_{\{0, \dots, i\}} \quad (j \in [k(P)]). \quad (13)$$

Proposition 7.8. Let $R = (R_i)_{i \in \mathbb{N}}$ be a compatible sequence of recursive blow-ups relative to $N = (N_i)_{i \in \mathbb{N}}$ and let $V = (V_i)_{i \in \mathbb{N}}$ be given by $V_i \stackrel{\text{def}}{=} V(N_i)$.

Then R is convergent and the following hold for its limit $\phi_R \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$.

i. If R is conservative, then $\phi_R = \phi_{\mathcal{N}^N}$.

ii. There exists a $T_{\mathcal{L}}$ -on \mathcal{H} over Ω^V with $\phi_R = \phi_{\mathcal{H}}$

$$\mathcal{N}_P^N \subseteq \mathcal{H}_P \subseteq \mathcal{E}_{k(P)}(\Omega^V) \setminus \mathcal{N}_P^{\bar{N}} \text{ a.e.}$$

for every $P \in \mathcal{L}$, where $\bar{N} = (\bar{N}_i)_{i \in \mathbb{N}}$ is the sequence of complementary canonical \mathcal{L} -structures given by

$$V(\bar{N}_i) \stackrel{\text{def}}{=} V(N_i), \quad P^{\bar{N}_i} \stackrel{\text{def}}{=} (V(N_i))_{k(P)} \setminus P^{N_i} \quad (P \in \mathcal{L}).$$

iii. If $P(N)$ is the set of structures M such that there exist infinitely many $i \in \mathbb{N}$ with $M \cong N_i$, then $P(N) \subseteq P(\phi_R)$.

iv. If $\prod_{i \in \mathbb{N}} (1 - 1/|N_i|) = 0$, then for every positive measure $A \subseteq \Omega^V$, there exists $i \in \mathbb{N}$ such that $t_{\text{ind}}(N_i, \mathcal{H}|_A^F) > 0$ for every measure-isomorphism F modulo 0 from Ω_A^V to Ω^V .

Proof. To show that R is convergent, for each $i \in \mathbb{N}$, define the Euclidean structure \mathcal{N}^i in \mathcal{L} over Ω^V by

$$\mathcal{N}_P^i \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{k(P)}(\Omega^V) \mid R_i \vDash P(t_i^P(x))\},$$

where $t_i^P(x)$ is given by (13), that is, \mathcal{N}^i is the natural ‘‘step’’ Euclidean structure associated with R_i over Ω^V .

First note that since R is compatible, for every $i, j \in \mathbb{N}$, we have

$$\begin{aligned} & \sum_{P \in \mathcal{L}} \nu^V(\mathcal{N}_P^i \triangle \mathcal{N}_P^{i+j}) \\ & \leq \sum_{P \in \mathcal{L}} \nu^V(\{x \in \mathcal{E}_{k(P)}(\Omega^V) \mid \exists a, b \in [k(P)], (a \neq b \wedge x_{\{a\}}|_{\{0, \dots, i\}} = x_{\{b\}}|_{\{0, \dots, i\}})\}) \\ & \leq \sum_{P \in \mathcal{L}} \binom{k(P)}{2} \cdot \sum_{\sigma \in \prod_{\ell=0}^i V_\ell} \nu^V(K_{\sigma, V})^2 \\ & = \sum_{P \in \mathcal{L}} \binom{k(P)}{2} \cdot \prod_{\ell=0}^i |V_i|^{-1} \xrightarrow{i \rightarrow \infty} 0. \end{aligned} \quad (14)$$

Therefore, it follows that for every finite \mathcal{L} -structure K , the limit $\lim_{i \rightarrow \infty} t_{\text{ind}}(K, \mathcal{N}^i)$ exists.

On the other hand, it is also straightforward to check that for every finite \mathcal{L} -structure K , we have

$$||R_i|^{|\mathcal{K}|} \cdot t_{\text{ind}}(K, \mathcal{N}^i) - |T_{\text{ind}}(K, R_i)|| \leq |R_i|^{|\mathcal{K}|} - (|R_i|)_{|\mathcal{K}|} \leq O_K(|R_i|^{|\mathcal{K}|-1})$$

hence we get

$$\lim_{i \rightarrow \infty} t_{\text{ind}}(K, R_i) = \lim_{i \rightarrow \infty} t_{\text{ind}}(K, \mathcal{N}^i),$$

that is, $R = (R_i)_{i \in \mathbb{N}}$ is convergent.

Consider now the case when R is conservative. Then the same argument used in (14) gives

$$\sum_{P \in \mathcal{L}} \nu^V(\mathcal{N}_P^i \Delta \mathcal{N}_P^N) \leq \sum_{P \in \mathcal{L}} \binom{k(P)}{2} \cdot \prod_{\ell=0}^i |V_i|^{-1} \xrightarrow{i \rightarrow \infty} 0,$$

so item (i) follows.

To prove item (ii), note that (14) implies that for each $P \in \mathcal{L}$, the sequence of indicator functions $(\mathbb{1}_{\mathcal{N}_P^i})_{i \in \mathbb{N}}$ is convergent in $L^1(\mathcal{E}_{k(P)}(\Omega^V))$, so let f_P be their L^1 -limit. Since f_P is also the a.e. limit of $(\mathbb{1}_{\mathcal{N}_P^i})_{i \in \mathbb{N}}$, it must be a.e. $\{0, 1\}$ -valued, so there exists \mathcal{H}_P such that $f_P = \mathbb{1}_{\mathcal{H}_P}$ a.e. Finally, L^1 -convergence implies that $\lim_{i \rightarrow \infty} t_{\text{ind}}(K, \mathcal{N}^i) = t_{\text{ind}}(K, \mathcal{H})$.

We claim that for every $x \in \mathcal{N}_P^N$, there exists $i_0 \in \mathbb{N}$ such that $x \in \mathcal{N}_P^i$ for every $i \geq i_0$. Indeed, if $x \in \mathcal{N}_P^N$, then there exists $i_0 \in \mathbb{N}$ such that $R^{(N_0, \dots, N_{i_0})} \models P(t_{i_0}^P(x))$. The definition of the conservative recursive blow-ups $R^{(N_0, \dots, N_i)}$ implies that $R^{(N_0, \dots, N_i)} \models P(t_i^P(x))$ for every $i \geq i_0$. From the minimality of the conservative recursive blow-ups, we get $R^i \models P(t_i^P(x))$, hence $x \in \mathcal{N}_P^i$ for every $i \geq i_0$. Since $\mathbb{1}_{\mathcal{N}_P^i}$ converges a.e. to $\mathbb{1}_{\mathcal{H}_P}$, we conclude that $\mathcal{N}_P^N \subseteq \mathcal{H}_P$ a.e.

By a symmetric argument, it follows that for every $x \in \mathcal{N}_P^{\bar{N}}$, there exists $i_0 \in \mathbb{N}$ such that $x \in \mathcal{N}_P^i$ for every $i \geq i_0$, from which we conclude that $\mathcal{N}_P^{\bar{N}} \subseteq \mathcal{E}_{k(P)}(\Omega^V) \setminus \mathcal{H}_P$ a.e. and thus $\mathcal{H}_P \subseteq \mathcal{E}_{k(P)}(\Omega^V) \setminus \mathcal{N}_P^{\bar{N}}$ a.e.

Let us now show item (iii). Fix $M \in P(N)$ and let us show that $M \in P(\phi_R)$. By Lemma 7.2, it is sufficient to show that for every positive measure $A \subseteq \Omega^V$ and every measure-isomorphism F modulo 0 from Ω_A^V to Ω^V , we have $M \in Q(\mathcal{H}_A^F)$.

Let then $\epsilon > 0$ be such that $\epsilon < 1/|M|$. By Lemma 4.11, there exists $t_0 \in \mathbb{N}$ such that for every $t \geq t_0$, there exists $\sigma \in \prod_{\ell=0}^{t-1} V_\ell$ such that $\nu^V(A \cap K_{\sigma, V}) \geq (1 - \epsilon) \cdot \nu^V(K_{\sigma, V})$. Since $M \in P(N)$, there exists $t \geq t_0$ such that $M \cong N_t$. Since $\{K_{(\sigma, u), V} \mid u \in V_t\}$ partitions $K_{\sigma, V}$ into $|V_t| = |M|$ parts of equal measure, it follows that for every $u \in V_t$, we have

$$\nu^V(A \cap K_{(\sigma, u), V}) \geq \left(1 - \epsilon - \frac{|M| - 1}{|M|}\right) \cdot \nu^V(K_{\sigma, V}) > 0.$$

Note now that if $x \in \mathcal{E}_{V_t}(\Omega^V)$ is such that $x_{\{u\}} \in A \cap K_{(\sigma, u), V}$ for every $u \in V_t$, then $x \in T_{\text{ind}}(N_t, \mathcal{H}|_A^F)$. Thus $t_{\text{ind}}(N_t, \mathcal{H}|_A^F) > 0$, hence $M = N_t \in Q(\mathcal{H}|_A^F)$, as desired.

It remains to show item (iv). Suppose not, that is, suppose that there exist some positive measure $A \subseteq \Omega^V$ and some measure-isomorphism F modulo 0 from Ω_A^V to Ω^V such that for every $i \in \mathbb{N}$, we have $t_{\text{ind}}(N_i, \mathcal{H}|_A^F) = 0$.

Let $n \in \mathbb{N}$ be large enough so that $\prod_{i=0}^{n-1} (1 - 1/|N_i|) < \nu^V(A)$ and let

$$\Sigma \stackrel{\text{def}}{=} \left\{ \sigma \in \prod_{i=0}^{n-1} V_i \mid \nu^V(A \cap K_{\sigma, V}) > 0 \right\}.$$

We claim that for every $m \in \{0, \dots, n-1\}$ and every $\tau \in \prod_{i=1}^{m-1} V_i$, there exists $u_\tau \in V_m$ such that (τ, u_τ) is not a prefix of any element of Σ . Suppose not, that is, suppose that there exist $m \in \{0, \dots, n-1\}$ and $\tau \in \prod_{i=1}^{m-1} V_i$ such that for every $u \in V_m$, there exists some $\sigma^u \in \Sigma$ such that (τ, u) is a prefix of σ^u . But then the set of $x \in \mathcal{E}_{V_m}(\Omega^V)$ such that $x_{\{u\}} \in A \cap K_{\sigma^u, V}$ for every $u \in V_m$ is a positive measure set that is contained in $T_{\text{ind}}(N_m, \mathcal{H}|_A^F)$, contradicting the fact that $t_{\text{ind}}(N_m, \mathcal{H}|_A^F) = 0$. Thus the claim is proved.

Let now Σ^* be the set of $\sigma \in \prod_{i=0}^{n-1} V_i$ such that for every $m \in \{0, \dots, n-2\}$, we have $u_{\sigma|_{\{0, \dots, m-1\}}} \neq \sigma_m$. Our last claim shows that $\Sigma \subseteq \Sigma^*$. Now it is easy to see that

$$\nu^V(A) = \sum_{\sigma \in \Sigma} \nu^V(A \cap K_{\sigma, V}) \leq \sum_{\sigma \in \Sigma^*} \nu^V(K_{\sigma, V}) = \prod_{i=0}^{n-1} \left(1 - \frac{1}{|N_i|}\right) < \nu^V(A),$$

a contradiction. Thus, item (iv) is proved. \square

Theorem 7.9. *The following are equivalent for a family \mathcal{F} of finite \mathcal{L} -structures (up to isomorphism) containing at least one structure of size at least 2.*

i. *The family \mathcal{F} is strongly persistent.*

ii. *The family \mathcal{F} is weakly closed under substitutions and closed under substructures.*

Furthermore, if all predicate symbols of \mathcal{L} have arity at most 2, then the above are also equivalent to:

iii. *The family \mathcal{F} is persistent.*

Proof. The implication (i) \implies (ii) follows from Lemma 7.3(ii) as $\mathcal{F} = P(\phi)$ for some *weakly random* $\phi \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$.

For the implication (ii) \implies (i), let \mathcal{F}' be the set of elements of \mathcal{F} of size at least 2 and let $N = (N_i)_{i \in \mathbb{N}}$ be an enumeration of all elements of \mathcal{F}' that repeats each element of \mathcal{F}' infinitely often. Since \mathcal{F} is weakly closed under substitutions and closed substructures, by Remark 3.10, it follows that $\mathcal{F} = \mathcal{F}' \cup \{K_0, F_1\}$ for some \mathcal{L} -structure F_1 of size 1 (and where K_0 is the trivial \mathcal{L} -structure of size 0).

By Lemma 7.6, there exists a compatible sequence $R = (R_i)_{i \in \mathbb{N}}$ of recursive blow-ups relative to N with $R_i \in \mathcal{F}$ for every $i \in \mathbb{N}$ and by Proposition 7.8(iii), we know that R converges to some $\phi_R \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ such that $\mathcal{F}' = P(N) \subseteq P(\phi_R)$ and since $P(\phi_R)$ is closed under substructures (see Lemma 7.3) and $\mathcal{F} = \mathcal{F}' \cup \{K_0, F_1\}$, we must have $\mathcal{F} \subseteq P(\phi_R)$. On the other hand, since $R_i \in \mathcal{F}$, it follows that $P(\phi_R) \subseteq Q(\phi_R) \subseteq \mathcal{F}$, hence $\mathcal{F} = Q(\phi_R) = P(\phi_R)$ as desired.

If all predicate symbols of \mathcal{L} have arity at most 2, then implication (iii) \implies (ii) follows from Lemma 7.3(i) as $\mathcal{F} = P(\phi)$ for some $\phi \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ and the implication (i) \implies (iii) is obvious. \square

Again, the assumption of arity at most 2 is crucial for the inclusion of item (iii) in the equivalence of Theorem 7.9 as illustrated by Example 7.4.

We conclude this section by observing operations that preserve the notions discussed so far. The next proposition shows naturality of the operators Q and P and of the weak randomness property in the sense that the operators P and Q commute with open interpretations and weak randomness is preserved by open interpretations.

Proposition 7.10. *Let $I: T_1 \rightsquigarrow T_2$ be an open interpretation. The following hold for $\phi \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$.*

- i. *We have $Q(\phi^I) = I(Q(\phi))$.*
- ii. *We have $P(\phi^I) = I(P(\phi))$.*
- iii. *If ϕ is weakly random, then so is ϕ^I .*

Proof. Item (i) follows directly from the definition of ϕ^I , see (6).

Item (ii) follows directly from item (i) and the fact that if ψ is a sub-object of ϕ , then ψ^I is a sub-object of ϕ^I and conversely, every sub-object of ϕ^I is of the form ψ^I for some sub-object ψ of ϕ .

Item (iii) follows trivially from items (i) and (ii). \square

Before we proceed, we recall the notion of couplings and independent couplings of limits from [CR23, Definitions 2.3, 2.4 and 2.5], which played a key role in the study of the natural quasirandomness properties **UCouple**[ℓ] and **UInduce**[ℓ] in that work.

Definition 7.11. Given canonical theories T_1 and T_2 in finite relational languages \mathcal{L}_1 and \mathcal{L}_2 , respectively, the *disjoint union* $T_1 \cup T_2$ is the canonical theory in the disjoint union language $\mathcal{L}_1 \dot{\cup} \mathcal{L}_2$ whose axioms are those of T_1 (about predicate symbols in \mathcal{L}_1) and those of T_2 (about predicate symbols in \mathcal{L}_2), that is, the models of $T_1 \cup T_2$ correspond to a model of T_1 and a model of T_2 on the same vertex set.

A *coupling* of $\phi_1 \in \text{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and $\phi_2 \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ is a limit $\psi \in \text{Hom}^+(\mathcal{A}[T_1 \cup T_2], \mathbb{R})$ such that $\phi_i = \psi^{I_i}$ for every $i \in [2]$, where $I_i: T_i \rightsquigarrow T_1 \cup T_2$ is the *structure-erasing interpretation* that acts identically on predicate symbols of T_i .

The *independent coupling* of $\phi_1 \in \text{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and $\phi_2 \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ is the limit $\phi_1 \otimes \phi_2 \in \text{Hom}^+(\mathcal{A}[T_1 \cup T_2], \mathbb{R})$ given by

$$(\phi_1 \otimes \phi_2)(M) \stackrel{\text{def}}{=} \frac{|\text{Aut}(M_1)| \cdot |\text{Aut}(M_2)|}{|M|! \cdot |\text{Aut}(M)|} \cdot \phi_1(M_1) \cdot \phi_2(M_2),$$

where $M_i \stackrel{\text{def}}{=} I_i(M)$. Alternatively, if \mathcal{N}^i ($i \in [2]$) is a T_i -on over Ω_i with $\phi_{\mathcal{N}^i} = \phi_i$, then we have $\phi_1 \otimes \phi_2 = \phi_{\mathcal{N}^1 \otimes \mathcal{N}^2}$ for the $(T_1 \cup T_2)$ -on $\mathcal{N}^1 \otimes \mathcal{N}^2$ over the product space $\Omega_1 \otimes \Omega_2$ given by

$$(\mathcal{N}^1 \otimes \mathcal{N}^2)_P \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{k(P)}(\Omega_1 \otimes \Omega_2) \mid \pi_{i,k(P)}(x) \in \mathcal{N}_P^i\}$$

whenever $P \in \mathcal{L}_i$ ($i \in [2]$), where $\pi_{i,k(P)}: \mathcal{E}_{k(P)}(\Omega_1 \otimes \Omega_2) \rightarrow \mathcal{E}_{k(P)}(\Omega_i)$ is the natural projection.

The next proposition says that weak randomness is preserved under independent couplings.

Proposition 7.12. *If $\phi_1 \in \text{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and $\phi_2 \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ are weakly random, then so is their independent coupling $\phi_1 \otimes \phi_2$.*

Proof. Let \mathcal{N}^i be a T_i -on over $\Omega_i = (X_i, \mathcal{A}_i, \mu_i)$ such that $\phi_i = \phi_{\mathcal{N}^i}$ and let $\Omega \stackrel{\text{def}}{=} \Omega_1 \otimes \Omega_2$. It is clear from the definition of $\phi_1 \otimes \phi_2$ that for every $M \in \mathcal{M}[T_1 \cup T_2]$, we have $M \in Q(\phi_1 \otimes \phi_2)$ if and only if $I_1(M) \in Q(\phi_1)$ and $I_2(M) \in Q(\phi_2)$, where $I_i: T_i \rightsquigarrow T_1 \cup T_2$ ($i \in [2]$) is the structure-erasing interpretation.

By Lemma 7.2(i), to show that $\phi_1 \otimes \phi_2$ is weakly random, it is sufficient to show that for every positive measure $A \subseteq \Omega$ and every measure-isomorphism F modulo 0 from Ω_A to Ω , we have $Q(\phi_1 \otimes \phi_2) = Q((\mathcal{N}^1 \otimes \mathcal{N}^2)|_A^F)$.

Let $M \in Q(\phi_1 \otimes \phi_2)$ and let us show that $M \in Q((\mathcal{N}^1 \otimes \mathcal{N}^2)|_A^F)$. For each $i \in [2]$, let $M_i \stackrel{\text{def}}{=} I_i(M)$ and let

$$B \stackrel{\text{def}}{=} \{(x, y) \in \mathcal{E}_{V(M)}(\Omega_1) \times X_2 \mid x \in T_{\text{ind}}(M_1, \mathcal{N}^1) \wedge \forall v \in V(M), (x_{\{v\}}, y) \in A\}.$$

Our objective is to show that $(\mu_1 \otimes \mu_2)(B) > 0$. To do so, for each $y \in X_2$, let

$$A(y) \stackrel{\text{def}}{=} \{x \in X_1 \mid (x, y) \in B\}$$

and note that Fubini's Theorem implies that the set

$$\tilde{X}_2 \stackrel{\text{def}}{=} \{y \in X_2 \mid \mu_1(A(y)) > 0\}$$

has positive μ_2 -measure.

Since ϕ_1 is weakly random, for every $y \in \tilde{X}_2$ and every measure isomorphism \tilde{F}_y modulo 0 from $(\Omega_1)_{A(y)}$ to Ω_1 , we have $t_{\text{ind}}(M_1, \mathcal{N}^1|_{A(y)}^{\tilde{F}_y}) > 0$, thus Fubini's Theorem gives

$$(\mu_1 \otimes \mu_2)(B) \geq \int_{\tilde{X}_2} t_{\text{ind}}(M_1, \mathcal{N}^1|_{A(y)}^{\tilde{F}_y}) \cdot \mu_1(A(y))^{|M|} d\mu_2(y) > 0.$$

For every $x \in T_{\text{ind}}(M_1, \mathcal{N}^1) \subseteq \mathcal{E}_{V(M)}(\Omega_1)$, define the set

$$B(x) \stackrel{\text{def}}{=} \{y \in X_2 \mid (x, y) \in B\} = \{y \in X_2 \mid \forall v \in V(M), (x_{\{v\}}, y) \in A\}$$

and note that Fubini's Theorem again implies that the set

$$\widetilde{T}_{\text{ind}}(M_1, \mathcal{N}^1) \stackrel{\text{def}}{=} \{x \in T_{\text{ind}}(M_1, \mathcal{N}^1) \mid \mu_2(B(x)) > 0\}$$

has positive μ_1 -measure. Since ϕ_2 is weakly random, for every $x \in \widetilde{T}_{\text{ind}}(M_1, \mathcal{N}^1)$ and every measure isomorphism \tilde{G}_x modulo 0 from $(\Omega_2)_{B(x)}$ to Ω_2 , we have $t_{\text{ind}}(M_2, \mathcal{N}^2|_{B(x)}^{\tilde{G}_x}) > 0$, thus Fubini's Theorem gives

$$t_{\text{ind}}(M, (\mathcal{N}^1 \otimes \mathcal{N}^2)|_A^F) \geq \int_{\widetilde{T}_{\text{ind}}(M_1, \mathcal{N}^1)} t_{\text{ind}}(M_2, \mathcal{N}^2|_{B(x)}^{\tilde{G}_x}) \cdot \mu_2(B(x))^{|M|} d\mu_1(x) > 0,$$

concluding the proof. \square

Remark 7.13. As we mentioned before, weak randomness can be seen as a weakening of the natural quasirandomness property **UIInduce**[1] of [CR23]. Since **UIInduce**[1] (and more generally, **UIInduce**[ℓ]) is not preserved under independent couplings, one can consider the class **UIInduce** $_{\otimes}[\ell]$ that is the closure of **UIInduce**[ℓ] under independent couplings and open interpretations and in [CR23, §10], it was asked if any of these classes yields a meaningful notion of randomness or if they are already “too large”. It was already noted in [CR23] that the quasirandom permuto (see Proposition 7.15) is in **UIInduce** $_{\otimes}[\ell]$ for every $\ell \in \mathbb{N}_+$ and that even the largest class **UIInduce** $_{\otimes}[1]$ among the **UIInduce** $_{\otimes}[\ell]$ does not contain all limits.

Since **UIInduce**[1] implies weak randomness, from Propositions 7.10(iii) and 7.12 it follows that every element of **UIInduce** $_{\otimes}[1]$ is weakly random; this further justifies the adjective “weak” in weak randomness: it is a quasirandomness notion weaker than the weakening **UIInduce** $_{\otimes}[1]$ of **UIInduce**[1] that is still meaningful.

Let us point out that there are weakly random limits that are not in **UIInduce** $_{\otimes}[1]$: namely, one can show that if W is a universal weakly random $\{0, 1\}$ -valued graphon of T_{Graph} (e.g., $\phi_W = \phi_G^*$ as in Lemma 4.13 for an enumeration $G = (G_m)_{m \in \mathbb{N}}$ of all finite graphs of size at least 2), then ϕ is weakly random but is not in **UIInduce** $_{\otimes}[1]$. However, since the length of the proof outweighs its enlightenment value, we omit it.

Recall from Definition 2.8 that a trivial limit $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is any limit of the form $\phi = \phi_{\mathcal{N}}$ for some theon \mathcal{N} whose peons all have measure in $\{0, 1\}$. For general couplings, the next proposition says that the coupling of a trivial limit with a weakly random limit is weakly random.

Proposition 7.14. *If ψ is a coupling of a trivial $\phi_1 \in \text{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and a weakly random $\phi_2 \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$, then ψ is weakly random.*

Proof. Let \mathcal{L}_1 and \mathcal{L}_2 be the languages of T_1 and T_2 , respectively.

Since ϕ_1 is trivial, it follows that \mathcal{L}_1 can be partitioned into $\mathcal{L}_1 = \mathcal{L}_1^0 \cup \mathcal{L}_1^1$ so that for every $M_1 \in Q(\phi_1)$ every $P \in \mathcal{L}_1$, we have

$$P^{M_1} = \begin{cases} \emptyset, & \text{if } P \in \mathcal{L}_1^0, \\ (V(M_1))_{k(P)}, & \text{if } P \in \mathcal{L}_1^1. \end{cases}$$

This implies that if ξ is a coupling of some $\zeta \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ with ϕ_1 , then

$$Q(\xi) = \{\widehat{M}_2 \mid M_2 \in Q(\zeta)\}, \quad (15)$$

where $\widehat{M}_2 \in \mathcal{M}_{V(M_2)}[T_1 \cup T_2]$ is given by

$$P^{\widehat{M}_2} \stackrel{\text{def}}{=} \begin{cases} P^{M_2}, & \text{if } P \in \mathcal{L}_2, \\ \emptyset, & \text{if } P \in \mathcal{L}_1^0, \\ (V(M_2))_{k(P)}, & \text{if } P \in \mathcal{L}_1^1. \end{cases}$$

Now since ϕ_1 is trivial, ϕ_1 is the only sub-object of ϕ_1 , which means that every sub-object ψ' of ψ is a coupling of ϕ_1 with some sub-object ϕ'_2 of ϕ_2 . Since ϕ_2 is weakly random, we have $Q(\phi'_2) = Q(\phi_2)$, hence $Q(\psi') = Q(\psi)$ follows since the right-hand side of (15) is the same for $(\xi, \zeta) = (\psi, \phi_2)$ and $(\xi, \zeta) = (\psi', \phi'_2)$. Therefore ψ is weakly random. \square

As a simple application of Propositions 7.10 and 7.12 above, let us prove Proposition 5.11 that says that the graphon of agreements of the quasirandom permutoon (see Figure 6) is a universal weakly random limit of $T_{\text{PermGraph}}$ by showing that the quasirandom permutoon $\psi_{\text{QR}} \in \text{Hom}^+(\mathcal{A}[T_{\text{Perm}}], \mathbb{R})$ has the same property for T_{Perm} . We point the reader interested in the theories of limits of permutations and quasirandom permutations to [Coo04, KP13, CKN⁺20].

Recall that the quasirandom permutoon is given by $\psi_{\text{QR}} \stackrel{\text{def}}{=} \phi_{\mathcal{N}^{\text{QR}}}$, where \mathcal{N}^{QR} is the T_{Perm} -on over $[0, 1]^2$ given by

$$\mathcal{N}_{\prec_i}^{\text{QR}} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_2([0, 1]^2) \mid \pi_i(x_{\{1\}}) < \pi_i(x_{\{2\}})\} \quad (i \in [2]),$$

where $\pi_i: [0, 1]^2 \rightarrow [0, 1]$ is the projection onto the i th coordinate.

Proposition 7.15. *The quasirandom permutoon ψ_{QR} is a universal weakly random limit of T_{Perm} .*

Proof. It is straightforward to check that $Q(\psi_{\text{QR}}) = \mathcal{M}[T_{\text{Perm}}]$. On the other hand, ψ_{QR} is the independent coupling $\psi \otimes \psi$ of the unique limit $\psi \in \text{Hom}^+(\mathcal{A}[T_{\text{LinOrder}}], \mathbb{R})$ of the theory of (strict) linear orders with itself. Since ψ is obviously weakly random (as T_{LinOrder} is finitely categorical), by Proposition 7.12, it follows that ψ_{QR} is weakly random. \square

We can now derive Proposition 5.11 that says that the graphon of agreements of the quasirandom permutoon is universal weakly random limit of $T_{\text{PermGraph}}$ as an easy consequence.

Proof of Proposition 5.11. The graphon of agreements of the quasirandom permuton represents the limit $\phi \stackrel{\text{def}}{=} \psi_{\text{QR}}^I$ for the open interpretation $I: T_{\text{Graph}} \rightsquigarrow T_{\text{Perm}}$ given by

$$I(E)(x, y) \stackrel{\text{def}}{=} x \neq y \wedge (x \prec_1 y \leftrightarrow x \prec_2 y),$$

so ϕ is weakly random by Propositions 7.10(iii) and 7.15.

Finally, Proposition 7.10(i) implies $Q(\phi) = I(Q(\psi_{\text{QR}})) = I(T_{\text{Perm}}) = T_{\text{PermGraph}}$. \square

Remark 7.16. It is easy to see that the same permutations used in the proof of Proposition 5.12 can be used to show that $\mathcal{M}[T_{\text{Perm}}]$ is closed under substitutions but not primally almost finite, hence $T_{\text{Perm}} \notin \text{WR}$. However, let us point out that had we proved only the result for T_{Perm} , this would not have immediately implied Proposition 5.12 as primality is not necessarily preserved under open interpretations (even though closure under substitutions is).

8 What about weak randomness in general?

In this brief section we provide a partial generalization of Theorem 5.4 of Section 5 to universal theories in finite relational languages. For the easier direction, we will only be able to generalize Lemma 5.6 when all arities are at most 2 (Proposition 8.2) and even though the harder direction will generalize directly in Proposition 8.4 below, this naive generalization is essentially empty when all arities are at least 3, as in this case there are only finitely many prime structures (see Remark 3.13). It is not clear at this point what form a characterization of WR should take in the presence of higher arity predicates.

Definition 8.1. We say that a canonical theory T in a finite relational language has the *weakly random Erdős–Hajnal property* (abbreviated $T \in \text{WR}$) if every $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ has a weakly random sub-object.

Proposition 8.2. Let \mathcal{L} be a finite relational language whose predicate symbols have arity at most 2 and let T be a canonical theory in \mathcal{L} . If $\mathcal{M}[T]$ is monochromatically primally almost finite, then $T \in \text{WR}$.

Proof. We prove this by the contra-positive. Suppose $T \notin \text{WR}$ and let us show that the set \mathcal{P} of monochromatic prime models of T is not almost finite. By Lemma 3.23, it is sufficient to present a sequence $(F_n)_{n \in \mathbb{N}}$ of finite monochromatic prime models of T such that F_n is not a substructure of F_m whenever $n < m$.

Since $T \notin \text{WR}$, there must exist a limit $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ that does not contain any weakly random sub-object.

We now construct a sequence $(\phi_n)_{n \in \mathbb{N}}$ of sub-objects of ϕ and a sequence $(F_n)_{n \in \mathbb{N}}$ of finite prime models of T satisfying the following.

- i. For every $n \in \mathbb{N}$, ϕ_{n+1} is a sub-object of ϕ_n .
- ii. For every $n \in \mathbb{N}$, $F_n \in Q(\phi_n) \setminus Q(\phi_{n+1})$.

We construct these sequences inductively as follows.

1. We claim that there exists a sub-object ϕ_0 of ϕ such that there exists $M_1 \in \mathcal{M}_1[T]$ with $\phi_0(M_1) = 1$ (and thus all $M \in \mathcal{M}_1[T] \setminus \{M_1\}$ have $\phi_0(M) = 0$). Indeed, if $M_1 \in \mathcal{M}_1[T]$ is such that $\phi(M_1) > 0$ and \mathcal{N} is an Euclidean structure over Ω with $\phi_{\mathcal{N}} = \phi$, then $A \stackrel{\text{def}}{=} T_{\text{ind}}(M_1, \mathcal{N})$ is a positive measure set, so for any measure-isomorphism F modulo 0 from Ω_A to Ω , the sub-object $\phi_0 \stackrel{\text{def}}{=} \phi_{\mathcal{N}|_A^F}$ satisfies the desired property.
2. Given $\phi_n \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$, since ϕ_n is a sub-object of ϕ , we know that ϕ_n is not weakly random, so there exists $N_n \in Q(\phi_n) \setminus P(\phi_n)$. Let \mathcal{P}_n be the set of substructures of N_n that are prime. By Lemma 7.3(i), we know that $P(\phi_n)$ is *strongly* closed under substitutions and since $N_n \in S(\mathcal{P}_n)$, there must exist $F_n \in \mathcal{P}_n \setminus P(\phi_n)$ and since $Q(\phi_n)$ is closed under substructures, we get $F_n \in Q(\phi_n) \setminus P(\phi_n)$. From the definition of $P(\phi_n)$, it then follows that there exists a sub-object ϕ_{n+1} of ϕ_n (hence also of ϕ) such that $F_n \in Q(\phi_n) \setminus Q(\phi_{n+1})$.

Let now $n, m \in \mathbb{N}$ be such that $n < m$. By induction, we know that ϕ_m is a sub-object of ϕ_{n+1} , so $Q(\phi_m) \subseteq Q(\phi_{n+1})$, which in turn implies that $F_n \in Q(\phi_n) \setminus Q(\phi_m)$. Since $Q(\phi_m)$ is closed under substructures and $F_m \in Q(\phi_m)$, it follows that F_n is not a substructure of F_m .

Finally, since all ϕ_n are also sub-objects of ϕ_0 , we must have $Q(\phi_n) \cap \mathcal{M}_1[T] \subseteq Q(\phi_0) \cap \mathcal{M}_1[T] = \{M_1\}$. This implies that for every *unary* predicate symbol $P \in \mathcal{L}$ and every $n \in \mathbb{N}$, we have $M_1 \models \forall x, P(x)$ if and only if $F_n \models \forall x, P(x)$ (otherwise, we would have $Q(\phi_n) \cap \mathcal{M}_1[T] \neq \{M_1\}$). Thus the F_n are monochromatic. \square

Lemma 8.3. *Let $N = (N_i)_{i \in \mathbb{N}}$ be a sequence of prime \mathcal{L} -structures of size at least 2 such that for each $i \in \mathbb{N}$, there are finitely many $j \in \mathbb{N}$ such that N_i is a substructure of N_j , let $R = (R_i)_{i \in \mathbb{N}}$ be a compatible sequence of recursive blow-ups relative to N and let $\phi_R \in \text{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ be the limit of R . If $\prod_{i \in \mathbb{N}} (1 - 1/|N_i|) = 0$, then ϕ_R does not have any weakly random sub-object.*

Proof. Let V and \mathcal{H} be as in Proposition 7.8. Suppose toward a contradiction that ϕ_R has a weakly random sub-object. By Lemma 7.2, there exists a positive measure set $A \subseteq \Omega^V$ and a measure-isomorphism F modulo 0 from Ω_A^V to Ω^V such that $\mathcal{H}|_A^F$ is weakly random. By Proposition 7.8(iv), there exists $i_0 \in \mathbb{N}$ such that $t_{\text{ind}}(N_{i_0}, \mathcal{H}|_A^F) > 0$.

Let $j_0 \stackrel{\text{def}}{=} \max\{j \mid N_{i_0} \text{ is a substructure of } N_j\} < \infty$. Since $\{K_{\sigma, V} \mid \sigma \in \prod_{\ell=0}^{j_0} V_{\ell}\}$ partitions Ω^V , there must exist $\sigma \in \prod_{\ell=0}^{j_0} V_{\ell}$ such that $\nu^V(A \cap K_{\sigma, V}) > 0$.

From the definition of \mathcal{H} , it follows that for every measure-isomorphism \tilde{F} modulo 0 from $\Omega_{K_{\sigma, V}}^V$ to Ω^V , we have $\phi_{\mathcal{H}|_{K_{\sigma, V}}^{\tilde{F}}} = \phi_{R'}$ for the sequence $R' = (R'_i)_{i \in \mathbb{N}}$ given by $R'_i \stackrel{\text{def}}{=} R_{j_0+1+i}|_{U_i}$, where

$$U_i \stackrel{\text{def}}{=} \left\{ \tau \in \prod_{\ell=0}^{j_0+1+i} V_{\ell} \mid \tau|_{\{0, \dots, j_0\}} = \sigma \right\}.$$

Note also that R' is a compatible sequence of recursive blow-ups relative to the shifted sequence $N' = (N'_i)_{i \in \mathbb{N}}$ given by $N'_i \stackrel{\text{def}}{=} N_{j_0+1+i}$.

We claim now that $t_{\text{ind}}(N_{i_0}, \mathcal{H}|_{K_{\sigma,V}}^{\tilde{F}}) = 0$. Indeed, since N_{i_0} is prime, for this density to be positive, N_{i_0} must be a substructure of infinitely many R'_i , but since $R'_i \in S(\{N_j \mid j \geq j_0+1\})$, Lemma 3.17 says that this can only happen if N_{i_0} is a substructure of some N_j with $j \geq j_0+1$, which would contradict the definition of j_0 .

Finally, this is a contradiction since $\mathcal{H}|_A^F$ was assumed to be weakly random but $N_{i_0} \in Q(\mathcal{H}|_A^F) \setminus P(\mathcal{H}|_A^F)$ as $\nu^V(A \cap K_{\sigma,V}) > 0$ and $t_{\text{ind}}(N_{i_0}, \mathcal{H}|_{K_{\sigma,V}}^{\tilde{F}}) = 0$. \square

Proposition 8.4. *Let T be a canonical theory such that $\mathcal{M}[T]$ is weakly closed under substitutions. If $T \in \text{WR}$, then T is primally almost finite.*

Before we prove Proposition 8.4, let us note that it is completely trivial when all arities are at least 3 as in this case there are only finitely many prime structures by Remark 3.13.

Proof. We prove this by the contra-positive. Suppose $\{N'_i \mid i \in \mathbb{N}\}$ is an infinite antichain of prime models of T and without loss of generality, assume every N'_i has size at least 2 (as $\mathcal{M}_0[T] \cup \mathcal{M}_1[T]$ is finite).

For each $n \in \mathbb{N}$, let $r_n \in \mathbb{N}_+$ be large enough so that $(1 - 1/|N'_i|)^{r_n} \leq 1/2$ and for each $\ell \in \mathbb{N}$, let $N_\ell \stackrel{\text{def}}{=} N'_n$ for the unique $n \in \mathbb{N}$ such that $\sum_{m=0}^{n-1} r_m \leq \ell < \sum_{m=0}^n r_m$. Clearly, for each $\ell \in \mathbb{N}$, there exist exactly r_ℓ values of $t \in \mathbb{N}$ such that N_ℓ is a substructure of N_t . Note also that

$$\prod_{\ell \in \mathbb{N}} \left(1 - \frac{1}{|N_\ell|}\right) = \prod_{m \in \mathbb{N}} \left(1 - \frac{1}{|N'_m|}\right)^{r_m} \leq \prod_{m \in \mathbb{N}} \frac{1}{2} = 0.$$

Since $\mathcal{M}[T]$ is weakly closed under substitutions, by Lemma 7.6, there exists a compatible sequence $R = (R_\ell)_{\ell \in \mathbb{N}}$ of recursive blow-ups relative to $N = (N_\ell)_{\ell \in \mathbb{N}}$ with $R_\ell \in \mathcal{M}[T]$ for every $\ell \in \mathbb{N}$, and by Lemma 8.3, the limit $\phi_R \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ of R does not have any weakly random sub-object, hence $T \notin \text{WR}$. \square

Let us conclude this section by observing operations that preserve WR (at the level of theories). The next proposition shows naturality (at the level of theories) of WR , that is, it is preserved by open interpretations.

Proposition 8.5. *If $I: T_1 \rightsquigarrow T_2$ is an open interpretation and $T_2 \in \text{WR}$, then $I(T_2) \in \text{WR}$.*

Proof. Follows from Proposition 7.10(iii), the fact that every $\phi \in \text{Hom}^+(\mathcal{A}[I(T_2)], \mathbb{R})$ is of the form $\phi = \psi^I$ for some $\psi \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ and the fact that if ψ is a sub-object of ϕ , then ψ^I is a sub-object of ϕ^I and conversely, every sub-object of ϕ^I is of the form ψ^I for some sub-object ψ of ϕ . \square

It is easy to see that WR is not preserved under disjoint unions of theories (see Definition 7.11): the theory of linear orders T_{LinOrder} satisfies WR (as it is finitely categorical) but the theory of permutations $T_{\text{Perm}} = T_{\text{LinOrder}} \cup T_{\text{LinOrder}}$ does not satisfy WR (see Remark 7.16). However, the next proposition says that WR at least interacts well with disjoint unions with theories with AEHP (see Definition 2.8).

Proposition 8.6. *Let T_1 and T_2 be universal theories and suppose $T_1 \in \text{AEHP}$. Then the following hold.*

i. *If $T_2 \in \text{AEHP}$, then $T_1 \cup T_2 \in \text{AEHP}$.*

ii. *If $T_2 \in \text{WR}$, then $T_1 \cup T_2 \in \text{WR}$.*

To prove this proposition, we will need the following result from [CR23] on theons representing couplings (see Definition 7.11).

Proposition 8.7 ([CR23, Proposition 4.3]). *Let $\psi \in \text{Hom}^+(\mathcal{A}[T_1 \cup T_2], \mathbb{R})$ be a coupling of $\phi_1 \in \text{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and $\phi_2 \in \text{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ and let \mathcal{N}^1 be a T_1 -on over Ω such that $\phi_1 = \phi_{\mathcal{N}^1}$. Then there exists a $(T_1 \cup T_2)$ -on \mathcal{H} over $\Omega \otimes \Omega$ such that $\psi = \phi_{\mathcal{H}}$ and $\mathcal{H}_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)$ for every predicate symbol P in the language of T_1 (when we naturally identify $\mathcal{E}_{k(P)}(\Omega \otimes \Omega)$ with $\mathcal{E}_{k(P)}(\Omega) \times \mathcal{E}_{k(P)}(\Omega)$).*

Proof of Proposition 8.6. Let $\psi \in \text{Hom}^+(\mathcal{A}[T_1 \cup T_2], \mathbb{R})$ and let $I_i: T_i \rightsquigarrow T_1 \cup T_2$ ($i \in [2]$) be the structure-erasing interpretation. Then ψ is a coupling of $\phi_1 \stackrel{\text{def}}{=} \psi^{I_1}$ and $\phi_2 \stackrel{\text{def}}{=} \psi^{I_2}$.

Let \mathcal{N}^1 be a T_1 -on over $\Omega = (X, \mathcal{A}, \mu)$ such that $\phi_1 = \phi_{\mathcal{N}^1}$. Since $T_1 \in \text{AEHP}$, by [CM22, Theorem 5.11], there exists a positive measure set $A \subseteq X$ and a measure-isomorphism F modulo 0 from Ω_A to Ω such that $\phi_{\mathcal{N}^1|_A^F}$ is trivial.

Let now \mathcal{H} be the $(T_1 \cup T_2)$ -on over $\Omega' \stackrel{\text{def}}{=} \Omega \otimes \Omega$ given by Proposition 8.7 and let $\mu' \stackrel{\text{def}}{=} \mu \otimes \mu$ be the underlying measure of Ω' . Let also $A' \stackrel{\text{def}}{=} A \times X$ and let $F' \stackrel{\text{def}}{=} F \otimes \text{id}_X$ be the measure-isomorphism modulo 0 from $\Omega'_{A'}$ to Ω' that acts as F on the first coordinate and acts identically on the second coordinate.

Suppose $T_2 \in \text{AEHP}$. Since $I_2(\mathcal{H}|_{A'}^{F'}) = I_2(\mathcal{H})|_{A'}^{F'}$ is a T_2 -on, by [CM22, Theorem 5.11], there exists a positive $\mu'_{A'}$ -measure set $B \subseteq X \times X$ such that $I_2(\mathcal{H})|_{A'}^{F'}|_B^{\tilde{F}}$ is trivial for every measure-isomorphism \tilde{F} modulo 0 from $(\Omega'_{A'})_B$ to $\Omega'_{A'}$.

Set $B' \stackrel{\text{def}}{=} B \cap A'$ so that B' is a positive μ' -measure set such that $I_2(\mathcal{H})|_{B'}^{\tilde{F} \circ F'}$ is trivial. Note now that since $\mathcal{H}_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)$ for every predicate symbol P in the language of T_1 , we get $\phi_{I_1(\mathcal{H})|_{A'}^{F'}} = \phi_{\mathcal{N}^1|_A^F}$, which is a trivial limit. Since $\phi_{I_1(\mathcal{H})|_{B'}^{\tilde{F} \circ F'}}$ is a sub-object of $\phi_{\mathcal{N}^1|_A^F}$, it must also be trivial. Hence, $\psi = \phi_{\mathcal{H}|_{B'}^{\tilde{F} \circ F'}}$ must be trivial as it is a coupling of two trivial limits $\phi_{\mathcal{H}|_{B'}^{\tilde{F} \circ F'}}^{I_1} = \phi_{I_1(\mathcal{H})|_{B'}^{\tilde{F} \circ F'}}$ and $\phi_{\mathcal{H}|_{B'}^{\tilde{F} \circ F'}}^{I_2} = \phi_{I_2(\mathcal{H})|_{B'}^{\tilde{F} \circ F'}}$. Thus item (i) is proved.

For item (ii), we make the same construction but taking $B \subseteq X \times X$ with positive $\mu'_{A'}$ -measure such that $I_2(\mathcal{H})|_{A'}^{F'}|_B^{\tilde{F}}$ is weakly random as guaranteed by $T_2 \in \text{WR}$. Then $\psi = \phi_{\mathcal{H}}$ must be weakly random by Proposition 7.14 as it is a coupling of a trivial limit $\phi_{\mathcal{H}|_{B'}^{\tilde{F} \circ F'}}^{I_1} = \phi_{I_1(\mathcal{H})|_{B'}^{\tilde{F} \circ F'}}$ with a weakly random limit $\phi_{\mathcal{H}|_{B'}^{\tilde{F} \circ F'}}^{I_2} = \phi_{I_2(\mathcal{H})|_{B'}^{\tilde{F} \circ F'}}$. \square

9 Conclusion and open problems

In this paper we studied the notion of weak randomness, a weakening of the quasirandomness property **UInduce**[1] (see [CR23]). In the language of graphs, a graphon is weakly random if

the set of finite graphs G having non-zero density is invariant across all subgraphons. In the more general language of structures, weak randomness requires the limit $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ to be such that for every sub-object ψ of ϕ and every finite structure M , we have $\phi(M) > 0$ if and only if $\psi(M) > 0$. We characterized (strongly) persistent families of structures, i.e., those that correspond to a theory T that has a universal weakly random limit (that is, a weakly random ϕ such that $\text{Th}(\phi) = T$) as precisely those that are closed under substructures and weakly closed under substitutions.

We also studied a weakening **WR** of **AEHP**. In the language of graphs, **WR** is a property of a hereditary class of graphs which requires that every graphon associated to the class contains a weakly random sub-graphon. We characterized **WR** for hereditary classes of graphs that are closed under substitution as precisely those classes which are “primally almost finite”, meaning that in the partial order on elements of the class given by induced subgraph, there is no infinite antichain of prime graphs. In the general language of structures, **WR** requires every limit of T to contain a weakly random sub-object (see Definitions 2.8, 5.1 and 8.1). We characterized **WR** for theories T with maximum arity at most 2 and $\mathcal{M}[T]$ closed under substitutions as precisely the set of theories T that are monochromatically primally almost finite.

A very natural open problem that was not addressed in this paper is to characterize weak randomness at the level of objects, that is, to provide an equivalent property to $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ being weakly random. Toward this goal, a natural first step is to ask how different can two weakly random objects ϕ and ψ be. A first source of difference is obviously that they can have different persistence sets $P(\phi) \neq P(\psi)$. On the other hand, if $P(\phi) = P(\psi)$, then we can attempt to measure their difference based on the sub-object partial pre-order and it is natural to ask what is the structure of the partially pre-ordered set $\Phi_{\mathcal{F}} \stackrel{\text{def}}{=} \{\phi \mid P(\phi) = Q(\phi) = \mathcal{F}\}$ for some (strongly) persistent class \mathcal{F} . Obviously, if $\mathcal{F} = \{K_n \mid n \in \mathbb{N}\}$ or $\mathcal{F} = \{\overline{K}_n \mid n \in \mathbb{N}\}$, then the set $\Phi_{\mathcal{F}}$ has only one element, but even for the next simplest case $\mathcal{F} = S(\{K_0, K_2, \overline{K}_2\})$ of induced subgraphs of recursive blow-ups of C_4 , the structure of the partial pre-order on $\Phi_{\mathcal{F}}$ is not clear: does it have incomparable elements? What about infinite antichains? By Proposition 7.8, if $G = (G_n)_{n \in \mathbb{N}}$ in which each G_n is either K_2 or \overline{K}_2 and both K_2 and \overline{K}_2 occur infinitely often, then the recursive blow-up ϕ_G satisfies $\phi_G \in \Phi_{\mathcal{F}}$ and we believe that changing the asymptotic proportion of edges and non-edges in G should produce incomparable elements of $\Phi_{\mathcal{F}}$.

As we mentioned in the introduction, the approximate Erdős–Hajnal property (**AEHP**) is a variation of the usual Erdős–Hajnal property (**EHP**) that allows for negligible errors, but requires linear-sized homogeneous sets in the presence of convergence. Since **WR** is a weakening of **AEHP**, we would like to ask the following more abstract question: what is the polynomial-sized error-free version of **WR** in the finite? Furthermore, since **AEHP** implies **EHP** and **WR** is a larger class than **AEHP**, is it still true that **WR** implies **EHP** for graphs? Of course, this implication must hold if the Erdős–Hajnal Conjecture is true. After the submission of this paper, Nguyen–Scott–Seymour posted a preprint with a proof of the Erdős–Hajnal conjecture for classes of graphs with bounded VC dimension [NSS24]. Combining their result with Theorem 6.10, one concludes that hereditary classes of graphs that are closed under

substitutions and satisfy **WR** must necessarily satisfy **EHP**, so it stands to reason to attempt to remove the closure under substitutions hypothesis.

As mentioned in Section 8 (see also Table 1), several of the proofs on weak randomness and the class **WR** do not generalize very well in the presence of predicates of arity at least 3. It is natural to ask if we can characterize **WR** in these cases in the presence of some simplifying assumption that would replace closure under substitution used in the binary case.

As briefly mentioned before, weak randomness is a weakening of the property **UIInduce**[1] of [CR23]. Since **UIInduce**[1] is part of a hierarchy of quasirandomness properties **UIInduce**[ℓ], one might expect that there exists a hierarchy of weak randomness as well. In turn, it may be that our difficulty in understanding **WR** in arity 3 comes from the fact that there is a wide variety of **UIInduce**[1] limits of 3-hypergraphs and since **UIInduce**[2] for 3-hypergraphons amounts again to only (full) quasirandom 3-hypergraphons, one might expect that the corresponding **WR**[2] property in arity 3 defined from an appropriate notion of “weak 2-randomness” (or more generally **WR**[$\ell - 1$] in arity ℓ) could be easier to handle. Since the definition of **UIInduce**[2] is considerably more technical than that of **UIInduce**[1] and our initial attempts at a weak 2-randomness definition did not yet yield any interesting results, we refrain from elaborating further.

Finally, in the absence of closure under substitutions, it is obvious that **WR** is no longer characterized by the primally almost finite condition: obvious counter-examples include the theories $T_{\omega \leq k}$ ($T_{\chi \leq k}$, resp.) of graphs whose clique number (chromatic number, resp.) is at most k , which clearly satisfy **AEHP** but are not closed under substitutions when $k \geq 2$ (as $K_{k+1} \in S(\{K_2\})$ is not a model of $T_{\omega \leq k}$ or $T_{\chi \leq k}$). It is possible to upgrade Lemma 5.6 and Proposition 8.2 to also cover the theories $T_{\omega \leq k}$ and $T_{\chi \leq k}$ via an interactive proof (more precisely, a two-player game in which the first player is attempting to show that some sub-object ψ must have $Q(\psi)$ monochromatically primally almost finite and the second player is attempting to deceive the first player), but we leave this result to a future work.

References

- [AC14] Ashwini Aroskar and James Cummings. Limits, regularity and removal for finite structures. Technical Report arXiv:1412.8084 [math.LO], arXiv e-print, 2014.
- [AHCH⁺18] Elad Aigner-Horev, David Conlon, Hiêp Hàn, Yury Person, and Mathias Schacht. Quasirandomness in hypergraphs. *Electron. J. Combin.*, 25(3):Paper 3.34, 22, 2018.
- [APS01] Noga Alon, János Pach, and József Solymosi. Ramsey-type theorems with forbidden subgraphs. volume 21, pages 155–170. 2001. Paul Erdős and his mathematics (Budapest, 1999).
- [Aus08] Tim Austin. On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probab. Surv.*, 5:80–145, 2008.

[CG90] F. R. K. Chung and R. L. Graham. Quasi-random hypergraphs. *Random Structures Algorithms*, 1(1):105–124, 1990.

[CGW89] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica*, 9(4):345–362, 1989.

[Chu90] Fan R. K. Chung. Quasi-random classes of hypergraphs. *Random Structures Algorithms*, 1(4):363–382, 1990.

[Chu14] Maria Chudnovsky. The Erdős-Hajnal conjecture—a survey. *J. Graph Theory*, 75(2):178–190, 2014.

[CK90] C. C. Chang and H. J. Keisler. *Model theory*, volume 73 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, third edition, 1990.

[CKN⁺20] Timothy F. N. Chan, Daniel Král’, Jonathan A. Noel, Yanitsa Pehova, Maryam Sharifzadeh, and Jan Volec. Characterization of quasirandom permutations by a pattern sum. *Random Structures Algorithms*, 57(4):920–939, 2020.

[CKOS16] Maria Chudnovsky, Ringi Kim, Sang-il Oum, and Paul Seymour. Unavoidable induced subgraphs in large graphs with no homogeneous sets. *J. Combin. Theory Ser. B*, 118:1–12, 2016.

[CM22] Leonardo Coregliano and Maryanthe Malliaris. Countable Ramsey. Technical Report arXiv:2203.10396 [math.CO], arXiv e-print, 2022.

[Coo04] Joshua N. Cooper. Quasirandom permutations. *J. Combin. Theory Ser. A*, 106(1):123–143, 2004.

[CR20] L. N. Coregliano and A. A. Razborov. Semantic limits of dense combinatorial objects. *Uspekhi Mat. Nauk*, 75(4(454)):45–152, 2020.

[CR23] L. N. Coregliano and A. A. Razborov. Natural quasirandomness properties. *Random Structures Algorithms*, 2023. To appear. ArXiv pre-print: arXiv:2012.11773 [math.CO].

[CRST06] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Ann. of Math. (2)*, 164(1):51–229, 2006.

[CS08] Maria Chudnovsky and Shmuel Safra. The Erdős-Hajnal conjecture for bull-free graphs. *J. Combin. Theory Ser. B*, 98(6):1301–1310, 2008.

[CSSS23] Maria Chudnovsky, Alex Scott, Paul Seymour, and Sophie Spirkl. Erdős-hajnal for graphs with no 5-hole. *Proceedings of the London Mathematical Society*, 126(3):997–1014, 2023.

[EH89] P. Erdős and A. Hajnal. Ramsey-type theorems. volume 25, pages 37–52. 1989. *Combinatorics and complexity (Chicago, IL, 1987)*.

[Gia97] Vassilis Giakoumakis. On the closure of graphs under substitution. *Discrete Math.*, 177(1-3):83–97, 1997.

[KP13] Daniel Král’ and Oleg Pikhurko. Quasirandom permutations are characterized by 4-point densities. *Geom. Funct. Anal.*, 23(2):570–579, 2013.

[KRS02] Yoshiharu Kohayakawa, Vojtěch Rödl, and Jozef Skokan. Hypergraphs, quasi-randomness, and conditions for regularity. *J. Combin. Theory Ser. A*, 97(2):307–352, 2002.

[LM15] John Lenz and Dhruv Mubayi. The poset of hypergraph quasirandomness. *Random Structures Algorithms*, 46(4):762–800, 2015.

[Lov12] László Lovász. *Large networks and graph limits*, volume 60 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2012.

[LS06] László Lovász and Balázs Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96(6):933–957, 2006.

[LS10] László Lovász and Balázs Szegedy. Regularity partitions and the topology of graphons. In *An irregular mind*, volume 21 of *Bolyai Soc. Math. Stud.*, pages 415–446. János Bolyai Math. Soc., Budapest, 2010.

[NSS24] Tung Nguyen, Alex Scott, and Paul Seymour. Induced subgraph density. VI. bounded VC-dimension. Technical Report arXiv:2312.15572 [math.CO], arXiv e-print, 2024.

[Pet13] F. Petrov. General removal lemma. Technical Report arXiv:1309.3795 [math.CO], arXiv e-print, 2013.

[Raz07] Alexander A. Razborov. Flag algebras. *J. Symbolic Logic*, 72(4):1239–1282, 2007.

[Sau72] N. Sauer. On the density of families of sets. *J. Combinatorial Theory Ser. A*, 13:145–147, 1972.

[She72] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J. Math.*, 41:247–261, 1972.

[SS97] Miklós Simonovits and Vera T. Sós. Hereditarily extended properties, quasi-random graphs and not necessarily induced subgraphs. *Combinatorica*, 17(4):577–596, 1997.

- [SS03] Miklós Simonovits and Vera T. Sós. Hereditary extended properties, quasi-random graphs and induced subgraphs. *Combinatorics, Probability and Computing*, 12(3):319–344, 2003.
- [Tho87] Andrew Thomason. Pseudorandom graphs. In *Random graphs '85 (Poznań, 1985)*, volume 144 of *North-Holland Math. Stud.*, pages 307–331. North-Holland, Amsterdam, 1987.
- [Tow17] Henry Towsner. σ -algebras for quasirandom hypergraphs. *Random Structures Algorithms*, 50(1):114–139, 2017.
- [VČ71] V. N. Vapnik and A. Ja. Červonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. *Teor. Verojatnost. i Primenen.*, 16:264–279, 1971.
- [Zve03] Igor Zverovich. Extension of hereditary classes with substitutions. *Discrete Appl. Math.*, 128(2-3):487–509, 2003.